

ONLINE LEARNING AS STOCHASTIC APPROXIMATION OF REGULARIZATION PATHS: OPTIMALITY AND ALMOST-SURE CONVERGENCE

PIERRE TARRÈS AND YUAN YAO

ABSTRACT. In this paper, an online learning algorithm is proposed as sequential stochastic approximation of a regularization path converging to the regression function in reproducing kernel Hilbert spaces (RKHSs). We show that it is possible to produce the best known strong (RKHS norm) convergence rate of batch learning, through a careful choice of the gain or step size sequences, depending on regularity assumptions on the regression function. The corresponding weak (mean square distance) convergence rate is optimal in the sense that it reaches the minimax and individual lower rates in the literature. In both cases we deduce almost sure convergence, using Bernstein-type inequalities for martingales in Hilbert spaces.

To achieve this we develop a bias-variance decomposition similar to the batch learning setting; the bias consists in the approximation and drift errors along the regularization path, which display the same rates of convergence, and the variance arises from the sample error analysed as a reverse martingale difference sequence. The rates above are obtained by an optimal trade-off between the bias and the variance.

1. INTRODUCTION

Consider the following classical problem of learning from examples: given a sequence of i.i.d. random examples $(z_t = (x_t, y_t))_{t \in \mathbb{N}}$ drawn from a probability measure ρ on $\mathcal{X} \times \mathcal{Y}$, one seeks to approximate the *regression function*

$$f_\rho(x) := \int_{\mathcal{Y}} y d\rho_{\mathcal{Y}|x},$$

i.e. the conditional expectation of y given x . Recall that f_ρ minimizes the following mean square error

$$(1) \quad \mathcal{E}(f) = \int_{\mathcal{X} \times \mathcal{Y}} (f(x) - y)^2 d\rho.$$

Date: January 17, 2013.

2000 Mathematics Subject Classification. 62L20, 68Q32, 68T05.

Key words and phrases. Online Learning, Stochastic Approximations, Regularization Path, Reproducing Kernel Hilbert Space.

The research of P.T. was supported in part by a Leverhulme Prize. The research of Y.Y. was supported in part by National Basic Research Program of China (973 Program 2011CB809105, 2012CB825501), NSFC grant 61071157, NSF grant 0325113, and a professorship in the Hundred Talents Program at Peking University. Part of this work was done while both authors visited Toyota Technological Institute at Chicago in the spring of 2005, written in an earlier version in [Yao 2006; Tarrès and Yao 2008].

The error of the approximation f of f_ρ is estimated for instance through the norm $\|f - f_\rho\|_\infty$ or $\|f - f_\rho\|_\rho$, where

$$\|f\|_\rho = \|f\|_{\mathcal{L}^2_{\rho_X}} = \left\{ \int_X |f(x)|^2 d\rho_X \right\}^{1/2},$$

(ρ_X being the marginal distribution of ρ on X), or through other norms in Hilbert spaces which, as we shall see later, may capture different regularity features of this approximation.

An *online learning algorithm* aims at obtaining this approximation of the regression function recursively, using at each time step the new example $z_t = (x_t, y_t)$ to update the current hypothesis f_{t-1} (approximating f_ρ) to f_t . In other words, $f_t = T_t(f_{t-1}, z_t)$ for some map $T_t : \mathcal{H} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{H}$, where \mathcal{H} is a Hilbert space of functions from \mathcal{X} to \mathcal{Y} , see for example [Smale and Yao 2006].

On the contrary, *batch learning* algorithms process a sample set given once and for all at some fixed time m , *i.e.* $\mathbf{z} = \{(x_i, y_i)\}_{i=1}^m$. The classical bias-variance paradigm is that of a trade-off between the requirement to fit the data, *i.e.* to provide a small empirical error

$$\hat{\mathcal{E}}_{\mathbf{z}}(f) := \frac{1}{m} \sum_{i=1}^m (f(x_i) - y_i)^2,$$

and the size of the space in which f can take place, in order to limit the impact of the noise created by the data. For instance, a Tikhonov regularization (or Ridge Regression) procedure as in [Evgeniou, Pontil, Poggio 2000] yields, given $\lambda > 0$,

$$f_{\mathbf{z}, \lambda} := \arg \min_{f \in \mathcal{H}} \left\{ \frac{1}{m} \sum_{i=1}^m (f(x_i) - y_i)^2 + \lambda \|f\|_{\mathcal{H}}^2 \right\}.$$

For more background on regularization of inverse problems, see for instance [Engl, Hanke, and Neubauer 1996]. In modern statistics, an L_1 -type regularization called LASSO [Tibshirani 1996], is proposed in pursuit of sparsity of f_ρ with respect to certain basis.

The regularization parameter λ is chosen as a function of the sample size m , and of some prior knowledge on the regularity of the function f_ρ . In this setting, probabilistic upper bounds of $\|f_m - f_\rho\|_{\mathcal{H}}$ were obtained for instance in [Cucker and Smale 2002, Smale and Zhou 2005].

In online learning, the sample size t is changing over time, so that the regularization parameter needs to be updated at each time step, and follows the *regularization path* defined as follows. Let, for all $\lambda > 0$, f_λ be the solution of the regularized least square problem

$$(2) \quad f_\lambda = \arg \min_{f \in \mathcal{H}} \mathcal{E}(f) + \lambda \|f\|_{\mathcal{H}}^2.$$

Depending on assumptions on the Hilbert space \mathcal{H} and on the regularity of f_ρ , f_λ converges to f_ρ in $\mathcal{L}^2_{\rho_X}$ or \mathcal{H} -norm when $\lambda \rightarrow 0$. The map

$$\begin{aligned} f_\cdot : \mathbb{R}_+ &\longrightarrow \mathcal{H} \\ \lambda &\longmapsto f_\lambda \end{aligned}$$

is called *regularization path* of f_ρ in \mathcal{H} .

Regularization paths gain rising attention from statistics, in particular in the LASSO case [Efron, Johnstone, Hastie, and Tibshirani 2004], where they are piecewise linear with respect to the parameter, which enables one to track the entire path with nearly the same amount of computational cost as a single fixed regularization. This property generalizes to the case where the loss

and the penalty are respectively piecewise quadratic and linear [Rosset and Zhu 2007]; note that this however does not include Tikhonov regularization.

Our purpose in this paper is to iteratively define an “online” sequence of functions $(f_t)_{t \in \mathbb{N}} \in \mathcal{H}$, which will provide a stochastic approximation of the Tikhonov regularization path $(f_{\lambda_t})_{t \in \mathbb{N}} \in \mathcal{H}$. With an adequate choice of the regularization parameters $\lambda_t \rightarrow 0$ based on a bias-variance trade-off, we show such a sequential stochastic approximation to be *optimal* in the sense that it reaches minimax and individual lower bound rates of convergence.

Our algorithms can be regarded as stochastic gradient descent algorithms to solve (2) with time varying regularization parameter λ_t , an extension from early works [Smale and Yao 2006; Yao 2010] which investigate the convergence $f_t \rightarrow f_\lambda$ for fixed $\lambda_t = \lambda > 0$. In that case a weak probabilistic upper bound for $\|f_t - f_\lambda\|$ was first proposed in [Smale and Yao 2006], based on Markov’s inequality. Improved upper bounds were later obtained in [Yao 2010], leading in some cases to the same rate of convergence of $(f_t)_{t \in \mathbb{N}}$ to f_λ as in batch learning given t examples.

However, as we shall see in this paper, time-varying λ_t was not addressed so far and leads to a more complicated bias-variance decomposition, whose heuristics is related to the existence of a phase transition in the convergence rate in stochastic approximation. We refer the reader to [Duflo 1996; Kushner and Yin 2003] for background on stochastic algorithms.

As in previous studies, we choose in this paper the Hilbert space \mathcal{H} to be a reproducing kernel Hilbert space (RKHS) \mathcal{H}_K for some kernel K . RKHS enables one to analyze nonparametric regressions in a coordinate-free manner, and the gradient descent method then takes an especially simple form [e.g. Kivinen, Smola, and Williamson 2004]. Moreover, RKHS provides a unified framework in several important settings, e.g.

- (i) generalized smooth spline functions in Sobolev spaces [Wahba 1990],
- (ii) real analytic functions with bounded bandwidth [Daubechies 1992] and their generalizations [Smale and Zhou 2004],
- (iii) gaussian processes [Loéve 1948; Parzen 1961].

In fact, any Hilbert space of functions on \mathcal{X} with a bounded evaluation functional is a RKHS [Wahba 1990]. By choosing suitable kernels, \mathcal{H}_K can be used to approximate any function in $\mathcal{L}_{\rho_{\mathcal{X}}}^2$; see for instance [Berlinet and Thomas-Agnan 2004] for wider background on RKHS.

Our analysis starts in the setting of a general Hilbert space \mathcal{W} in Section 3, with the study of an iteratively defined sequence, which is a stochastic approximation of the solution of some linear equation. This study will be specialized in later sections to the cases $\mathcal{W} = \mathcal{H}_K$ or $\mathcal{L}_{\rho_{\mathcal{X}}}^2$ in order to show the main results of the paper. Two structural decomposition theorems are introduced in that Section 3, the *reversed martingale decomposition* and the *martingale decomposition*, and play an important role in the proof of the main results, the former being suitable for strong convergence in \mathcal{H}_K and the latter for weak convergence in $\mathcal{L}_{\rho_{\mathcal{X}}}^2$.

Both decompositions lead to the breakdown of the total error $f_t - f_\rho$ into four parts: the *initial error* caused by the initial guess f_0 , the *sample error* as a reverse martingale difference sequence, the *approximation error* $f_{\lambda_t} - f_\rho$, and the *drift error* along the regularization path (f_{λ_t}) caused by time-varying λ_t .

By a suitable choice of step sizes, the initial error won't affect the convergence rates. Now a key observation is that the drift error, which does not appear in previous fixed regularization settings with $\lambda_t = \lambda$, has the same order as the approximation error. Bernstein-type inequalities for martingales in Hilbert spaces are then used to bound the sample error. Therefore we have a similar bias-variance decomposition in online learning as in batch learning, with the bias being the approximation and the drift errors, and the variance being the sample error. It is then possible to optimize in order to yield the same optimal rates in online learning as in batch learning.

The main theorems in this paper provide some probabilistic upper bounds for the convergence of $(f_t)_{t \in \mathbb{N}}$ to f_ρ , in \mathcal{H}_K or $\mathcal{L}_{\rho_{\mathcal{X}}}^2$, under the assumption that $f_\rho \in \mathcal{H}_K$ has additional regularity. The convergence rate in $\mathcal{L}_{\rho_{\mathcal{X}}}^2$ is optimal in the sense that it reaches the minimax and individual lower rate. The convergence rate in \mathcal{H}_K meets the same best rates as in batch learning [Smale and Zhou 2005]. Both upper bounds depend on a logarithmic power $\alpha > 0$ of the confidence threshold δ (i.e. $O(\log^\alpha 1/\delta)$). They imply by Borel-Cantelli Lemma the almost sure convergence of f_t to f_ρ in \mathcal{H}_K and $\mathcal{L}_{\rho_{\mathcal{X}}}^2$. Such a theorem improves on our early result in 2006 (see [Yao 2006]), where in mean square distance the upper bounds depended polynomially on the confidence (i.e. $O(\delta^{-\alpha})$), and whence solves the open problem raised therein.

The paper is organized as follows. Section 2 collects the main results. Section 3 studies stochastic approximations of regularization paths for linear operator equations in general Hilbert spaces, where the key martingale and reverse martingale decompositions are presented. Section 4 collects some estimates on the drift along the regularization path, $\|f_\lambda - f_\mu\|$ ($\lambda, \mu > 0$), which are needed for the study of the bias, i.e. the approximation and drift errors. Sections 5 and 6 respectively yield upper bounds for convergence in \mathcal{H}_K and $\mathcal{L}_{\rho_{\mathcal{X}}}^2$. Appendix A derives a probabilistic inequality from the Pinelis-Bernstein inequality for martingales in Hilbert spaces, which is used to derive the probabilistic upper bounds in this paper. Appendix B collects some preliminary upper bounds used in the paper. Appendix C gives proofs of some results in Section 3.2.

2. MAIN RESULTS

2.1. Notations and Assumptions. Let $\mathcal{X} \subseteq \mathbb{R}^n$ be closed, $\mathcal{Y} = \mathbb{R}$ and $\mathcal{L} = \mathcal{X} \times \mathcal{Y}$. Let ρ be a probability measure on \mathcal{L} , $\rho_{\mathcal{X}}$ be the induced marginal probability measure on \mathcal{X} , and let $\rho_{\mathcal{Y}|x}$ be the conditional probability measure on \mathcal{Y} with respect to $x \in X$. Define $f_\rho : \mathcal{X} \rightarrow \mathcal{Y}$ by $f_\rho(x) = \int_{\mathcal{Y}} y d\rho_{\mathcal{Y}|x}$, the *regression function* of ρ . In the sequel, we let $\mathbb{E}[\cdot]$ be the expectation with respect to ρ .

Let $\mathcal{L}_{\rho_{\mathcal{X}}}^2$ be the Hilbert space of square integrable functions with respect to $\rho_{\mathcal{X}}$. In the sequel $\|\cdot\|_\rho$ denotes the norm in $\mathcal{L}_{\rho_{\mathcal{X}}}^2$.

Let $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a *Mercer kernel*, i.e. a continuous¹ symmetric real function which is *positive semi-definite* in the sense that $\sum_{i,j=1}^m c_i c_j K(x_i, x_j) \geq 0$ for any $m \in \mathbb{N}$ and any choice of $x_i \in X$ and $c_i \in \mathbb{R}$ ($i = 1, \dots, m$). A Mercer kernel K induces a function $K_x : \mathcal{X} \rightarrow \mathbb{R}$ ($x \in \mathcal{X}$) defined by $K_x(x') = K(x, x')$. Let \mathcal{H}_K be the *reproducing kernel Hilbert space* (RKHS) associated with a Mercer kernel K , i.e. the completion of the $\text{span}\{K_x : x \in \mathcal{X}\}$ with respect to the inner

¹In computer science literature, one often bears in mind some implicit feature map $\Phi : \mathcal{X} \rightarrow \mathcal{H}$ which takes an input vector x to a high (or infinite) dimensional feature vector, say an element of a Hilbert space \mathcal{H} , and then one considers explicitly the inner product $K(x, x') = \langle \Phi(x), \Phi(x') \rangle$ as the kernel. In this construction, the continuity of K is equivalent to continuity of the feature map Φ .

product, defined as the linear extension of the bilinear form $\langle K_x, K_{x'} \rangle_K = K(x, x')$ ($x, x' \in \mathcal{X}$). The norm of \mathcal{H}_K is denoted by $\|\cdot\|_K$. The most important property of RKHS is the *reproducing property*: for all $f \in \mathcal{H}_K$ and $x \in X$, $f(x) = \langle f, K_x \rangle_K$.

Throughout this paper, assume that

Finiteness Condition. (A) There exists a constant $\kappa \geq 0$ such that

$$\kappa := \sup_{x \in \mathcal{X}} \sqrt{K(x, x)} < \infty.$$

(B) There exists a constant $M_\rho \geq 0$ such that

$$\text{supp}(\rho) \subseteq \mathcal{X} \times [-M_\rho, M_\rho].$$

Define the linear map

$$L_K : \mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{H}_K$$

by the following integral transform

$$L_K(f)(x) := \int_X K(x, t) f(t) d\rho_X(t).$$

It is well-known that L_K is well-defined, and that composition with the inclusion $\mathcal{H}_K \hookrightarrow \mathcal{L}_{\rho_{\mathcal{X}}}^2$ yields a compact positive self-adjoint operator on $\mathcal{L}_{\rho_{\mathcal{X}}}^2$ [e.g. [Halmos and Sunder 1978](#), [Cucker and Smale 2002](#)]. The restriction $L_K|_{\mathcal{H}_K} : \mathcal{H}_K \rightarrow \mathcal{H}_K$ is the *covariance operator* of $\rho_{\mathcal{X}}$ in \mathcal{H}_K : by the reproducing property,

$$L_K|_{\mathcal{H}_K} = \mathbb{E}[\langle \cdot, K_x \rangle_K].$$

Abusing notation, we will denote the three operators by L_K in the sequel.

Note that, by Cauchy-Schwarz inequality, $\|L_K f\|_\infty \leq \kappa^2 \|f\|_{\mathcal{L}_{\rho_{\mathcal{X}}}^2}$, so that

$$(3) \quad \|L_K\|_{\mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{L}_{\rho_{\mathcal{X}}}^2} \leq \kappa^2.$$

The compactness of $L_K : \mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{L}_{\rho_{\mathcal{X}}}^2$ implies the existence of an orthonormal eigensystem $(\mu_\alpha, \phi_\alpha)_{\alpha \in \mathbb{N}}$ in $\mathcal{L}_{\rho_{\mathcal{X}}}^2$. Recall that (see [Cucker and Smale 2002](#) for instance)

$$\sum_{\alpha \in \mathbb{N}} \mu_\alpha = \int_{\mathcal{X}} K(x, x) d\rho_{\mathcal{X}}(x) \leq \kappa^2.$$

We assume in this paper that all eigenvalues μ_α are positive. We can define, for all $r > 0$,

$$(4) \quad \begin{aligned} L_K^r : \quad \mathcal{L}_{\rho_{\mathcal{X}}}^2 &\rightarrow \mathcal{L}_{\rho_{\mathcal{X}}}^2 \\ \sum_{\alpha \in \mathbb{N}} a_\alpha \phi_\alpha &\mapsto \sum_{\alpha \in \mathbb{N}} a_\alpha \mu_\alpha^r \phi_\alpha; \end{aligned}$$

L_K^r can be regarded as a low-pass filter, and $\|L_K^r\| = \max_{\alpha \in \mathbb{N}} \mu_\alpha^r = \|L_K\|^r$. Note that $L_K^{1/2} : \mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{H}_K$ is an isometrical isomorphism of Hilbert spaces. Hence the eigenfunctions $(\phi_\alpha)_{\alpha \in \mathbb{N}}$ are orthogonal both in $\mathcal{L}_{\rho_{\mathcal{X}}}^2$ and \mathcal{H}_K .

For all $f \in \mathcal{L}_{\rho_{\mathcal{X}}}^2$ and $r > 0$, we write $L_K^{-r} f \in \mathcal{L}_{\rho_{\mathcal{X}}}^2$ when f lies in the image of the mapping $L_K^r : \mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{L}_{\rho_{\mathcal{X}}}^2$. Note that, if $r \geq 1/2$, then this implies $f \in \mathcal{H}_K$ because of the isometry $L_K^{1/2}$ between $\mathcal{L}_{\rho_{\mathcal{X}}}^2$ and \mathcal{H}_K .

For all $\lambda > 0$ and $r \in \mathbb{R} \setminus \{0\}$, we can similarly define $(L_K + \lambda I)^r : \mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{L}_{\rho, \mathcal{X}}^2$, which is a bijection, since $\sum_{\alpha \in \mathbb{N}} a_\alpha^2 < \infty$ is equivalent to $\sum_{\alpha \in \mathbb{N}} a_\alpha^2 (\lambda + \mu_\alpha)^r < \infty$, using $\mu_\alpha \rightarrow_{\alpha \rightarrow \infty} 0$.

It can be shown [e.g. Cucker and Smale 2002] that for any $\lambda \in \mathbb{R}_+$, the solution of (2) is

$$(5) \quad f_\lambda = (L_K + \lambda I)^{-1} L_K f_\rho \in \mathcal{H}_K.$$

In this paper, by $B_1, C_1, D_1, B_2, C_2, D_2, \dots$, we denote various constants, which are defined “locally” in the sense that the same notations appeared in different sections has different meanings.

2.2. Stochastic Gradient Algorithms. Let $\mathcal{F} = (\mathcal{F}_t)_{t \in \mathbb{N}_0} \in \mathcal{X} \times \mathcal{Y}$ be the filtration $\mathcal{F}_t = \sigma\{(x_i, y_i) : 1 \leq i \leq t\}$. In the sequel denote by $\mathbb{E}_t = \mathbb{E}[\cdot | \mathcal{F}_t]$, the conditional expectation w.r.t. \mathcal{F}_t . Consider the following \mathcal{F}_t -adapted process $(f_t)_{t \in \mathbb{N}}$ taking values in \mathcal{H}_K ,

$$(6) \quad f_t = f_{t-1} - \gamma_t [(f_{t-1}(x_t) - y_t) K_{x_t} + \lambda_t f_{t-1}], \quad \text{for some fixed } f_0 \in \mathcal{H}_K, \text{ e.g. } f_0 := 0$$

where

- (I) for each t , (x_t, y_t) is independent and identically distributed (i.i.d.) according to ρ ;
- (II) the *gain (step size) sequence* $(\gamma_t)_{t \in \mathbb{N}}$ and *regularization sequence* $(\lambda_t)_{t \in \mathbb{N}}$ are taking values in $\mathbb{R}_+ := (0, \infty)$, and converging to 0 as t goes to infinity.

Remark 2.1. The computational cost of this algorithm typically is $O(t^2)$. As each step t , the main computational cost is due to the evaluation $f_{t-1}(x_t)$ which needs to access all K_{x_i} ($1 \leq i \leq t$) in $O(t)$ steps. Thus the total cost is of $O(t^2)$ at time t . In the cases that one can store and access the values $f_t(x)$ for all x , e.g. on a grid of \mathcal{X} , the computational cost is merely linear $O(t)$ at the requirement of large memory and fast memory access.

By reproducing property, we can see that the gradient map of

$$V_z(f) = \frac{1}{2}[(f(x) - y)^2 + \lambda \|f\|_K^2], \quad z = (x, y) \in \mathcal{Z}$$

is given by $\text{grad}V_z(f) = (f(x) - y)K_x + \lambda f$ [e.g. Smale and Yao 2006], as a random variable depending on z . Since the expectation $\mathbb{E}[V_z(f)] = 2(\mathcal{E}(f) + \lambda \|f\|_K^2)$, algorithm (6) can thus be regarded as stochastic approximations of gradient descent method to solve (2), for each $\lambda = \lambda_t$.

2.3. Main Theorems. Theorem A provides sufficient conditions for the convergence of the *online learning* sequence $(f_t)_{t \in \mathbb{N}_0}$ in (6) to the regression function f_ρ . Theorem B and C explicit the corresponding convergence rates, respectively in \mathcal{H}_K and $\mathcal{L}_{\rho, \mathcal{X}}^2$.

Theorem A (Sufficient conditions for convergence). *Assume $f_\rho \in \mathcal{H}_K$, and let (f_t) be defined by equation (6), with assumptions (I)-(II). Then*

$$\limsup_{t \rightarrow \infty} \mathbb{E}[\|f_t - f_\rho\|_K^2] = 0,$$

if the following conditions are satisfied:

$$(A) \quad \sum_{t \rightarrow \infty} \gamma_t \lambda_t = \infty.$$

$$(B) \quad \limsup_{t \rightarrow \infty} \sum_{k=1}^n \gamma_k^2 \prod_{i=k+1}^n (1 - \gamma_i \lambda_i)^2 = 0,$$

$$(C) \limsup_{t \rightarrow \infty} \sum_{k=1}^n \|f_{\lambda_k} - f_{\lambda_{k-1}}\|_K \prod_{i=k+1}^n (1 - \gamma_i \lambda_i) = 0.$$

This theorem will be proved in Section 3, as a consequence of Theorem 3.5 in the setting of Hilbert spaces. Assumptions (B) and (C) can be replaced by the stronger (but less technical) assumptions (B') and (C') in Corollary 3.7 that $\gamma_t/\lambda_t \rightarrow 0$ and $\|f_{\lambda_t} - f_{\lambda_{t-1}}\|_K/(\lambda_t \gamma_t) \rightarrow 0$.

Remark 2.2. Although $\lambda_t \rightarrow 0$, condition (A) puts a restriction that $\gamma_t \lambda_t$ can not drop too fast, in fact this is necessary to “forget” the error caused by the initial guess f_0 . Condition (B) says that the step size $\gamma_t \rightarrow 0$, and it has to drop faster than the regularization parameter λ_t . Such a condition is to attenuate the random fluctuation caused by sampling. Condition (C) implies that the drifts of the regularization path (f_{λ_t}) converges to zero, at a speed faster than $\gamma_t \lambda_t$. This condition says that in the long run, the drifts along the regularization path should decrease fast enough for the algorithm to follow the path. The drifts depend on regularity of f_ρ , that the smoother f_ρ is, the faster drifts go down.

In the next two theorems (B) and (C) we choose the sequences $(\gamma_t)_{t \in \mathbb{N}}$ and $(\lambda_t)_{t \in \mathbb{N}}$ in order to optimize the rates of convergence in \mathcal{H}_K and $\mathcal{L}_{\rho, \mathcal{X}}^2$. This optimization is twofold.

First, the study of convergence of approximations of ordinary differential equations generically yields a phase transition between a slower rate with “shadowing” of mean-field trajectories, and a faster one, normally distributed after renormalization. Even though the picture is more complicated in our case, in particular because the vector f_t is infinite-dimensional, this justifies here that we choose $\gamma_t \lambda_t$ reciprocally linear in t .

Second, optimization over (γ_t) at fixed $(\gamma_t \lambda_t)$ yields a bias-variance trade-off similar to the one observed in statistical “batch” learning, which relies on the regularity assumption on the regression function f_ρ .

More precisely, let us first recall the phase transition in classical finite-dimensional stochastic approximation, in the rate of convergence towards a stable equilibrium. Naturally, we study the projections of the algorithm on the base of eigenvectors of the linearization of the ordinary differential equation at the equilibrium. Let $(\eta_t)_{t \in \mathbb{N}}$ be one of these projections, and assume for instance that the corresponding eigenvalue is -1 , so that the stochastic recursion is of the form

$$\eta_{t+1} = \eta_t + \gamma_t(-\eta_t + \epsilon_{t+1} + r_{t+1}),$$

where $\mathbb{E}_{t-1}[\epsilon_t] = 0$, (ϵ_t) is bounded, and (r_t) is small. For simplicity we will assume that $r_t = 0$ (which corresponds to the special case $\lambda_t = \lambda$ is a constant), but the heuristics holds on to the general case where r_t is less than quadratic in all coordinates. Let, for all $t \in \mathbb{N}$, $\beta_t := \prod_{k=1}^t (1 - \gamma_k)$. Then it is easy to show by induction that

$$\eta_t = \beta_t \left[\eta_0 + \sum_{j=1}^t \frac{\gamma_j}{\beta_j} \epsilon_j \right].$$

Now suppose for instance that $\gamma_t \sim c/t$ ($c > 0$); then $\beta_n n^c \xrightarrow[n \rightarrow \infty]{} C > 0$. Depending on the choice of c , η_t exhibits the following phase transition at $c = 1/2$ in its asymptotic dynamics.

- If $c < 1/2$ then $\sum(\gamma_j/\beta_j)^2 < \infty$, therefore $\sum_{j=1}^t \gamma_j \epsilon_j / \beta_j$ converges a.s. by Doob's convergence theorem, which implies that $\eta_t t^c \xrightarrow[t \rightarrow \infty]{} C'$ (where C' is a positive random variable), in other words that (η_t) asymptotically “shadows” one solution of the ODE

$$\frac{dx}{dt} = -cx.$$

- On the contrary, if $c > 1/2$, then $\sum(\gamma_j/\beta_j)^2 = \infty$, and by the martingale convergence theorem (see for instance [Williams 1991]), assuming for instance $\mathbb{E}_{t-1}[\epsilon_t^2] = D^2 > 0$ constant, and $\eta_t \sqrt{t}$ converges towards a centered normally distributed random variable with variance $a^2 D^2 / (2a - 1)$, and follows an associated Ornstein-Uhlenbeck process, see [Duflo 1996] for instance.

Therefore it suffices to choose $c > 1/2$ to achieve fast convergence rates. In this paper we will set $c = 1$ and choose $\gamma_t \lambda_t \sim 1/t$ to meet the heuristics above.

The next two theorems present some probabilistic upper bounds which characterize the convergence rates in \mathcal{H}_K and $\mathcal{L}_{\rho, \mathcal{X}}^2$, under certain regularity assumptions on the regression function f_ρ .

Let $t_0 > 0$ and, for all $t \in \mathbb{N}$,

$$\bar{t} := t + t_0,$$

where t_0 is large enough which won't affect the speed of convergence. We assume, in the statement of Theorems B and C, that, for all $t \in \mathbb{N}$,

$$\gamma_t = a \left(\frac{1}{\bar{t}} \right)^{\frac{2r}{2r+1}}, \quad \lambda_t = \frac{1}{a} \left(\frac{1}{\bar{t}} \right)^{\frac{1}{2r+1}}.$$

Theorem B (Upper Bounds for \mathcal{H}_K -convergence). *Assume $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r \in (1/2, 3/2]$, $a \geq 1$, and $t_0^\theta \geq a\kappa^2 + 1$. Then, for all $t \in \mathbb{N}$, with probability at least $1 - \delta$,*

$$\|f_t - f_\rho\|_K \leq \frac{C_0}{\bar{t}} + \left(C_1 a^{1/2-r} \log \frac{2}{\delta} + C_2 a \right) \left(\frac{1}{\bar{t}} \right)^{\frac{2r-1}{4r+2}},$$

where

$$C_0 := 2t_0^{\frac{4r+3}{4r+2}} M_\rho, \quad C_1 := \frac{20r-2}{(2r-1)(2r+3)} \|L_K^{-r} f_\rho\|_\rho, \quad C_2 := \frac{20(\kappa+1)^2 M_\rho}{\kappa}.$$

Its proof is given in Section 5.

Remark 2.3. Given $\delta > 0$, M_ρ and $\|L_K^{-r} f_\rho\|_\rho$, one can optimize a in order to minimize

$$h(a) := C_1 a^{1/2-r} \log \frac{2}{\delta} + C_2 a.$$

This yields the choice $a^* := [C_1 / ((r-1/2)C_2)]^{(r+1/2)^{-1}}$, with

$$h(a^*) = (r+1/2) \left[\frac{C_1 C_2^{r-1/2}}{(r-1/2)} \right]^{(r+1/2)^{-1}}.$$

This asymptotic rate in $M_\rho^{(2r-1)/(2r+1)} \|L_K^{-r} f_\rho\|_\rho^{2/(2r+1)} t^{-(r-1/2)/(4r+2)}$ is the same as the best known rates in batch learning algorithms; see [Theorem 2, Smale and Zhou 2005].

Remark 2.4. Note that the upper bound consists of three parts. The first term at a rate $O(t^{-1})$, captures the influence of the initial choice $f_0 = 0$, which does not depend on r and is faster than the remaining terms. The second term at a rate $O(\|L_K^{-r} f_\rho\|_\rho t^{-(2r-1)/(4r+2)})$, collects contributions from both drifts along the regularization path $f_{\lambda_t} - f_{\lambda_{t-1}}$ and the approximation error $f_{\lambda_t} - f_\rho$, since they share the same rates up to different constants. The third term at a rate $O(t^{-(2r-1)/(4r+2)})$, reflects the error caused by random fluctuations by the i.i.d. sampling. Later as we will see, the second term is a bound on the bias and the third term is a bound on the variance.

Theorem C (Upper Bounds for $\mathcal{L}_{\rho, \mathcal{X}}^2$ -convergence). *Assume that $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r \in [1/2, 1]$. Assume $a \geq 4$, and $t_0^\theta \geq 2 + 8\kappa^2 a$.*

Then, for all $t \in \mathbb{N}$, with probability at least $1 - \delta$ ($\delta \in (0, 1)$),

$$\|f_t - f_\rho\|_\rho \leq \frac{D_0}{t} + \left(D_1 a^{-r} + \sqrt{a} D_2 \log \frac{2}{\delta} \right) \left(\frac{1}{t} \right)^{\frac{r}{2r+1}} + \left(a^{3/2} D_3 \sqrt{\log t} + a^{5/2} D_4 \right) (\log(2/\delta))^2 \left(\frac{1}{t} \right)^{\frac{4r-1}{4r+2}}.$$

$$D_1 := \frac{5r+1}{r(1+r)} \|L_K^{-r} f_\rho\|_\rho, \quad D_2 := 10\kappa M_\rho, \quad D_3 = 63\kappa^2 M_\rho, \quad D_4 := 50\kappa^2 M_\rho t_0^{1/2-\theta}.$$

Its proof will be given in Section 6.

Remark 2.5. When $r \in (1/2, 1]$, the first term of $O(1/t)$ and the third term of $O(t^{-\frac{2r-1/2}{2r+1}} \log^{1/2} t)$ both drop faster than the second term of $O(t^{-\frac{r}{2r+1}})$, whence they can be ignored asymptotically. The second term as the dominant one, roughly speaking has contributions from two parts: the one with constant D_1 comes from the bias, *i.e.* the approximation and the drift errors, while the other with constant D_2 comes from the variance, *i.e.* the sample error.

Remark 2.6. A special case is $r = 1/2$, which is equivalent to say $f_\rho \in \mathcal{H}_K$. In this case $\gamma_t = \lambda_t = \bar{t}^{-1/2}$, whence it does not satisfy the Path Following Condition (B) in Theorem A. But Theorem C suggests a weaker notion that f_t follows the regularization path, *i.e.* $f_t \rightarrow f_\rho$ in $\mathcal{L}_{\rho, \mathcal{X}}^2$ rather than \mathcal{H}_K , which in fact converges at a rate of $O(t^{-1/4} \log^{1/2} t)$ uniformly for all $f_\rho \in \mathcal{H}_K$.

Remark 2.7. In all, the convergence in $\mathcal{L}_{\rho, \mathcal{X}}^2$ has rates $O(t^{-r/(2r+1)} \log^{1/2} t \cdot \log^2 1/\delta)$, a logarithmic polynomial on δ , whence the Borel-Cantelli Lemma implies almost sure convergence $\|f_t - f_\rho\|_{\mathcal{L}_{\rho, \mathcal{X}}^2} \xrightarrow{as} 0$. This result solves an open problem raised early in [Yao 2006].

Remark 2.8. To see the asymptotic optimality, consider the generalization error $\mathcal{E}(f) - \mathcal{E}(f_\rho) = \|f - f_\rho\|_\rho^2$ [e.g. see Cucker and Smale 2002]. Since the rate $O(t^{-r/(2r+1)})$ dominates when $r > 1/2$, then under the same condition of Theorem C, there holds with probability at least $1 - \delta$ ($\delta \in (0, 1)$), for all $t \in \mathbb{N}$,

$$\mathcal{E}(f_t) - \mathcal{E}(f_\rho) \leq O(t^{-2r/(2r+1)}).$$

For $r \in (1/2, 1]$, the asymptotic rate $O(t^{-2r/(2r+1)})$ has been shown to be optimal in the sense that it reaches the minimax and individual lower rate [Caponnetto and De Vito 2005]. To be precise, let $\mathcal{P}(b, r)$ ($b > 1$ and $r \in (1/2, 1]$) be the set of probability measure ρ on $\mathcal{X} \times \mathcal{Y}$, such that: (A) almost surely $|y| \leq M_\rho$; (B) $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$; (C) the eigenvalues $(\mu_n)_{n \in \mathbb{N}}$ of $L_K : \mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{L}_{\rho, \mathcal{X}}^2$, arranged in a nonincreasing order, are subject to the decay $\mu_n = O(n^{-b})$. Then the following minimax lower rate was given as Theorem 2 in [Caponnetto and De Vito 2005],

$$\liminf_{t \rightarrow \infty} \inf_{(z_i)_1^t \mapsto f_t} \sup_{\rho \in \mathcal{P}(b, r)} \mathbf{Prob} \left\{ (z_i)_1^t \in \mathcal{L}^t : \mathcal{E}(f_t) - \mathcal{E}(f_\rho) > C t^{-\frac{2rb}{2rb+1}} \right\} = 1$$

for some constant $C > 0$ independent on t , where the infimum in the middle is taken over all algorithms as a map $\mathcal{Z}^t \ni (z_i)_1^t \mapsto f_t \in \mathcal{H}_K$.

Note that in the minimax lower rate, the probability measure may change for different data size t , which violates the fundamental identical distribution assumption in learning. Therefore [Györfi, Kohler, Krzyżak, and Walk 2002] suggests a kind of individual lower rates for learning problems. The following individual lower rate was obtained as Theorem 3 in [Caponnetto and De Vito 2005]: for every $B > b$,

$$\inf_{((z_i)_1^t \mapsto f_t)_{t \in \mathbb{N}}} \sup_{\rho \in \mathcal{P}(b, r)} \limsup_{t \rightarrow \infty} \frac{\mathbb{E}[\mathcal{E}(f_t)] - \mathcal{E}(f_\rho)}{t^{-\frac{2rB}{2rB+1}}} > 0,$$

where the infimum is taken over arbitrary sequences of functions $f_t : \mathcal{Z}^t \rightarrow \mathcal{H}_K$. It can be seen that the key difference in the individual lower rate, lies in that by putting $\limsup_{t \rightarrow \infty}$ before $\sup_{\rho \in \mathcal{P}(b, r)}$, the probability measure ρ is applied to all sufficiently large t .

Now we compare these lower rates to our upper bound. Since $L_K : \mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{L}_{\rho, \mathcal{X}}^2$ is a trace-class operator, its eigenvalues are summable. Therefore by taking $b = B = 1$, one may obtain an eigenvalue-independent lower rate $O(t^{-2r/(2r+1)})$ for all possible L_K . Therefore, the upper bound by Theorem C reaches both the minimax and the individual lower rates.

3. SEQUENTIAL STOCHASTIC APPROXIMATIONS OF REGULARIZATION PATHS IN HILBERT SPACES

In this section, we study some stochastic approximation sequences in the more general setting of general Hilbert spaces.

Let \mathcal{W} be a Hilbert space with inner product $\langle \cdot, \cdot \rangle$ and associated norm $\|u\| := \sqrt{\langle u, u \rangle}$, and let $SL(\mathcal{W})$ be the vector space of self-adjoint bounded linear operators on \mathcal{W} , endowed with the canonical norm

$$\|A\| := \sup_{\|x\| \leq 1} \|Ax\|.$$

Let \mathcal{X} and \mathcal{Y} be two topological spaces (on which we make no other assumption), let $\mathcal{Z} := \mathcal{X} \times \mathcal{Y}$ and let ρ be a probability measure on the Borel σ -algebra of \mathcal{Z} . Let $A : \mathcal{Z} \rightarrow SL(\mathcal{W})$ and $b : \mathcal{Z} \rightarrow \mathcal{W}$ be random variables on the sample space \mathcal{Z} taking values respectively in $SL(\mathcal{W})$ and \mathcal{W} , and let

$$\bar{A} := \mathbb{E}[A], \quad \bar{b} := \mathbb{E}[b]$$

be their expectations on (\mathcal{Z}, ρ) .

Now assume that \bar{A} is a positive operator, hence invertible, but that it has an *unbounded* inverse. Knowing A and b , but not ρ (and subsequently not \bar{A} and \bar{b}), and assuming $\bar{b} \in \bar{A}(\mathcal{W})$, the aim is to devise a stochastic algorithm approximating the solution \bar{w} of the following linear equation

$$(7) \quad \bar{A}w = \bar{b},$$

using as data an i.i.d sequence $(z_t)_{t \in \mathbb{N}}$ in \mathcal{Z} with probability law ρ . As in the standard setting of Robbins-Monro (see [Robbins and Monro 1951], [Kiefer and Wolfowitz 1952]), it is natural to consider a stochastic gradient descent algorithm.

More precisely, the search for the solution \bar{w} of (7) is equivalent to the minimization of the quadratic potential map $\hat{V} : \mathcal{W} \rightarrow \mathbb{R}$

$$\hat{V}(w) := \frac{1}{2} \langle \bar{A}(w - \bar{w}), w - \bar{w} \rangle,$$

whose gradient $\text{grad } \hat{V} : \mathcal{W} \rightarrow \mathcal{W}$ is given by

$$\text{grad } \hat{V}(w) = \bar{A}w - \bar{b} = \mathbb{E}[Aw - b].$$

In the context of online learning presented in the first two sections, $\mathcal{W} := \mathcal{H}_K$, $A((x, y))(f) := f(x)K_x$, $b((x, y)) := yK_x$ (see Section 3.3), so that $\bar{A} = L_K$, $\bar{b} = L_K f_\rho$ and $\bar{w} = f_\rho$, and $\hat{V}(w) = \|f - f_\rho\|_{\mathcal{L}_{\rho, \mathcal{X}}^2}^2 = \mathcal{E}(f) - \mathcal{E}(f_\rho)$ is the generalization error.

A natural Robbins-Monro gradient descent algorithm would be

$$(8) \quad w_t = w_{t-1} - \gamma_t (A(z_t)w_{t-1} - b(z_t)),$$

since $\mathbb{E}_{z_t \sim \rho}[A(z_t)w_{t-1} - b(z_t)] = \bar{A}w_{t-1} - \bar{b}$.

However, the sample complexity analysis on Hilbert spaces, in order to estimate the sample size sufficient to approximate the minimizer with high probability, requires boundedness of \bar{A}^{-1} (see for instance [Smale and Yao 2006]).

To solve this ill-posed problem with unbounded \bar{A}^{-1} , one may construct sequences of random variables $(A_t)_{t \in \mathbb{N}}$ and $(b_t)_{t \in \mathbb{N}}$ on the sample space \mathcal{Z} taking values respectively in $SL(\mathcal{W})$ and \mathcal{W} , with the assumption that, if

$$\bar{A}_t := \mathbb{E}[A_t], \quad \bar{b}_t := \mathbb{E}[b_t]$$

are their expectations on (\mathcal{Z}, ρ) , \bar{A}_t has bounded inverse and $\bar{A}_t \rightarrow \bar{A}$, $\bar{b}_t \rightarrow \bar{b}$. Then the aim is to find assumptions ensuring that the stochastic approximation sequence $(w_t)_{t \in \mathbb{N}}$ iteratively defined by $w_0 := W_0$ deterministic, and

$$(9) \quad w_t = w_{t-1} - \gamma_t (A_t(z_t)w_{t-1} - b_t(z_t)),$$

where $(\gamma_t)_{t \in \mathbb{N}}$ is a real positive sequence, converges to the solution \bar{w} of (7) as t goes to infinity.

This question can be divided into two subquestions: first the *deterministic* convergence of

$$(10) \quad \bar{w}_t := \bar{A}_t^{-1} \bar{b}_t.$$

to \bar{w} , the path $t \mapsto \bar{w}_t$ being then called a *regularization path* of the solution of equation (7), and second the *probabilistic* convergence of the quantity

$$(11) \quad r_t := w_t - \bar{w}_t,$$

which we call the *remainder* (note that $\bar{w}_t = \bar{A}_t^{-1} \bar{b}_t \neq \mathbb{E}[w_t]$ in general). In the online learning case (see Section 3.3), we choose $A_t := A + \lambda_t I$, $(\lambda_t)_{t \in \mathbb{N}}$ positive sequence, $b_t := b$, so that $\bar{w}_t = f_{\lambda_t} \rightarrow f_\rho$ in \mathcal{H}_K .

We provide in Section 3.1 two structural decompositions of r_t , respectively a reversed martingale and a martingale one. Both expand r_t into three parts: one depending on the initial value of r , called the *initial error*, one depending on the *drift*

$$(12) \quad \Delta_j := \bar{w}_j - \bar{w}_{j-1}$$

along the regularization path (\bar{w}_t) called the *drift error*, and finally one random variable of zero mean called the *sample error*, respectively written as a reversed martingale and as a martingale at time t .

The reversed martingale decomposition will, on one hand, enable us to prove Theorem 3.5 below, whose corollary is Theorem A in the context of online learning, and which provides sufficient assumptions on the asymptotic behaviour of the norms of A_t , A_t^{-1} , \bar{A}_t^{-1} and b_t for the convergence of the variance of the remainder r_t . On the other hand, this reversed decomposition will yield Theorem B giving upper bounds on $f_t - f_\rho$ in \mathcal{H}_K with high probability, proved in Section 5.

The martingale decomposition will imply Theorem C giving upper bounds of $f_t - f_\rho$ in $\mathcal{L}_{\rho, \mathcal{X}}^2$ with high probability, proved in Section 6.

3.1. Two Structural Decomposition Theorems. For all $j, t \in \mathbb{N}$, let Π_j^t be the *random operator* on \mathcal{W} , on the sample space $\mathcal{L}^{\mathbb{N}}$, defined by

$$\Pi_j^t((z_i)_{i \in \mathbb{N}}) = \begin{cases} \prod_{i=j}^t (I - \gamma_i A_i(z_i)) & \text{if } j \leq t; \\ I & \text{otherwise.} \end{cases}$$

By a slight abuse of notation, we let $A_t := A_t(z_t)$ and $b_t := b_t(z_t)$ in the sequel, when there is no ambiguity.

Theorem 3.1 (Reversed Martingale Decomposition). *For all $s, t \in \mathbb{N}$, $t \geq s$,*

$$(13) \quad r_t = \Pi_{s+1}^t r_s - \sum_{j=s+1}^t \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j) - \sum_{j=s+1}^t \Pi_j^t \Delta_j$$

Remark 3.2. Note that Π_{j+1}^t is an operator whose randomness only depends on z_{j+1}, \dots, z_t , whereas the randomness in $A_j \bar{w}_j - b_j$, with zero mean, only depends on z_j . By independence of z_t , $t \in \mathbb{N}$, the conditional expectation $\mathbb{E}[\gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j) | z_{j+1}, \dots, z_t]$ is 0, whence for each t , $\gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j)$ is a *reversed martingale difference sequence* whose sum is a *reversed martingale sequence* with zero mean. For more background on reversed martingales, see for example [Neveu 1975].

Proof of Theorem 3.1. By definition,

$$\begin{aligned} r_t &= w_t - \bar{w}_t \\ &= w_{t-1} - \bar{w}_t - \gamma_t (A_t w_{t-1} - b_t) \\ &= (I - \gamma_t A_t) (w_{t-1} - \bar{w}_{t-1}) - (I - \gamma_t A_t) (\bar{w}_t - \bar{w}_{t-1}) + \gamma_t A_t (w_{t-1} - \bar{w}_t) - \gamma_t (A_t w_{t-1} - b_t) \\ &= (I - \gamma_t A_t) (w_{t-1} - \bar{w}_{t-1}) - (I - \gamma_t A_t) (\bar{w}_t - \bar{w}_{t-1}) - \gamma_t (A_t \bar{w}_t - b_t) \end{aligned}$$

which implies

$$(14) \quad r_t = (I - \gamma_t A_t) r_{t-1} - \gamma_t (A_t \bar{w}_t - b_t) - (I - \gamma_t A_t) \Delta_t.$$

The result follows by induction on $t \in \mathbb{N}$, $t \geq s$. \square

For all $j, t \in \mathbb{N}$, let

$$\chi_t = (\bar{A}_t - A_t) w_{t-1} + (b_t - \bar{b}_t),$$

and let $\bar{\Pi}_j^t$ be the *deterministic* operator on \mathcal{W} defined by

$$\bar{\Pi}_j^t = \begin{cases} \prod_{i=j}^t (I - \gamma_i \bar{A}_i) & \text{if } j \leq t; \\ I, & \text{otherwise.} \end{cases}$$

Theorem 3.3 (Martingale Decomposition). *For all $s, t \in \mathbb{N}$, $t \geq s$,*

$$(15) \quad r_t = \bar{\Pi}_{s+1}^t r_s + \sum_{j=s+1}^t \gamma_j \bar{\Pi}_{j+1}^t \chi_j - \sum_{j=s+1}^t \bar{\Pi}_j^t \Delta_j$$

Remark 3.4. The martingale decomposition was proposed in [Yao 2010]. Contrary to the reversed martingale decomposition, only the *sample error* is random here, the operator $\bar{\Pi}_{j+1}^t$ being deterministic. The process $(\gamma_j \bar{\Pi}_{j+1}^t \chi_j)_{j \in \mathbb{N}}$ is a *martingale difference sequence* since, for all $j \in \mathbb{N}$ and $t \geq j$, $\mathbb{E}_{j-1}[\gamma_j \bar{\Pi}_{j+1}^t \chi_j] = 0$. Note that the martingale property continues to hold for dependent sampling $z_t(z_1, \dots, z_{t-1})$, as long as $\mathbb{E}_{t-1}[A_t(z_t)] = \bar{A}_t$ and $\mathbb{E}_{t-1}[b_t(z_t)] = \bar{b}_t$.

The non-randomness of the operator $\bar{\Pi}_j^t$ will play a key role in the proof of Theorem C in the online learning context, since it will enable us to make explicit calculations involving the spectral decomposition of $L_K : \mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{L}_{\rho, \mathcal{X}}^2$ (recall that $\bar{A}_i = L_K + \lambda_i$ then). However, the fact that χ_t , contrary to $A_t \bar{w}_t - b_t$ in the reversed expansion, does not depend only on z_t but rather on the whole past $(z_i)_{0 \leq i \leq t}$, makes it necessary to obtain a preliminary upper bound of χ_t in Appendix C, which explains the factor $(\log 2/\delta)^2$ in Theorem C, rather than $\log 2/\delta$ in Theorem B.

Proof of Theorem 3.3. By definition,

$$\begin{aligned} r_t &= w_t - \bar{w}_t \\ &= w_{t-1} - \bar{w}_t - \gamma_t (A_t w_{t-1} - b_t) \\ &= (I - \gamma_t \bar{A}_t)(w_{t-1} - \bar{w}_t) + \gamma_t \chi_t, \text{ using } \bar{b}_t = \bar{A}_t \bar{w}_t \\ &= (I - \gamma_t \bar{A}_t)r_{t-1} + \gamma_t \chi_t - (I - \gamma_t \bar{A}_t)[\bar{w}_t - \bar{w}_{t-1}]. \end{aligned}$$

The result follows by induction on $t \in \mathbb{N}$, $t \geq s$. □

3.2. Sufficient Conditions for the Convergence of the Remainder. The following Theorem 3.5, which implies Theorem A in the context of online learning (see Section 3.3), states the convergence of $\|r_t\|^2 = \|w_t - \bar{w}_t\|^2$ to zero in expectation, under some assumptions on the asymptotic behaviour of the gain sequence γ_t and of the norms of b_t and operators A_t , A_t^{-1} and \bar{A}_t^{-1} .

The corresponding *Generalized Finiteness Condition* on the asymptotic behaviour of A_t and b_t is a generalization of the *Finiteness Condition* in [Smale and Yao 2006].

Generalized Finiteness Condition. Let $(\underline{\alpha}_t)_{t \in \mathbb{N}}$ and $(\bar{\alpha}_t)_{t \in \mathbb{N}}$ be deterministic positive sequences. For all $t \in \mathbb{N}$, assume that almost surely, A_t is positive, and the operators A_t , \bar{A}_t and \bar{A} are invertible (although \bar{A} has an *unbounded* inverse), and that

$$\|A_t\| \leq \bar{\alpha}_t, \quad \|A_t^{-1}\| \leq \underline{\alpha}_t^{-1}.$$

Theorem 3.5. *Consider the stochastic approximation sequence $(w_t)_{t \in \mathbb{N}_0}$ and remainder $(r_t)_{t \in \mathbb{N}_0}$ defined in (9)-(11).*

Suppose that the Generalized Finiteness Condition holds, and that the variance $\mathbb{E}\|A_t\bar{w}_t - b_t\|^2$ is uniformly bounded in $t \in \mathbb{N}$. Then

$$\mathbb{E}\|r_t\|^2 \rightarrow 0,$$

if the following assumptions hold:

$$(A) \gamma_t \rightarrow 0 \text{ and } \sum_t \gamma_t \underline{\alpha}_t = \infty,$$

$$(B) \limsup_{t \rightarrow \infty} \sum_{k=1}^n \gamma_k^2 \prod_{i=k+1}^n (1 - \gamma_i \underline{\alpha}_i)^2 = 0,$$

$$(C) \limsup_{t \rightarrow \infty} \sum_{k=1}^n \|\Delta_k\| \prod_{i=k+1}^n (1 - \gamma_i \underline{\alpha}_i) = 0.$$

The following Lemma 3.6 enables us to provide simple sufficient conditions for (B) and (C) in Corollary 3.7.

Lemma 3.6. *Let $(a_t)_{t \in \mathbb{N}}$ and $(b_t)_{t \in \mathbb{N}}$ be two real positive sequences converging to 0 when t goes to infinity. Then*

$$\limsup_{t \rightarrow \infty} a_t/b_t = 0 \text{ and } \sum_{t \in \mathbb{N}} b_t = \infty \implies \limsup_{t \rightarrow \infty} \sum_{k=1}^n a_k \prod_{i=k+1}^n (1 - b_i) = 0.$$

Corollary 3.7. *In the statement of Theorem 3.5, assumptions (B) and (C) may respectively be replaced by*

$$(B') \limsup_{t \rightarrow \infty} \frac{\gamma_t}{\underline{\alpha}_t} = 0,$$

$$(C') \limsup_{t \rightarrow \infty} \frac{\|\Delta_t\|}{\underline{\alpha}_t \gamma_t} = 0.$$

Theorem 3.5 and Lemma 3.6 are proved in Appendix C, and imply Corollary 3.7: Lemma 3.6 with $a_t := \gamma_t^2$ (resp. $a_t := \|\Delta_t\|$) and $b_t := \underline{\alpha}_t \gamma_t$ shows that (B) (resp. (C)) implies (B') (resp. (C')).

The proof of Theorem 3.5 makes use of the following preliminary Lemma 3.8 (shown in Appendix C), which implies some upper bounds of the norms of operators Π_j^t , $t \geq j$, also used in Sections 6 and 5.

Lemma 3.8. *Let $j_0 \in \mathbb{N}$, and let $(\gamma_t)_{t \in \mathbb{N}}$, $(\underline{\alpha}_t)_{t \in \mathbb{N}}$ and $(\bar{\alpha}_t)_{t \in \mathbb{N}}$ be real positive sequences, and let $(A_t)_{t \in \mathbb{N}}$ be a sequence of positive compact self-adjoint operators on the Hilbert space \mathcal{W} . Assume that, for all $t \geq j_0$, $\|A_t\| \leq \bar{\alpha}_t$, $\|A_t^{-1}\| \leq \underline{\alpha}_t^{-1}$ and $\gamma_t \bar{\alpha}_t \leq 1$.*

Then, for all $t \geq j_0$ and $j_0 \leq j \leq t$,

$$(A) \|I - \gamma_t A_t\| \leq 1 - \gamma_t \lambda_t;$$

$$(B) \|\Pi_j^t\| \leq \prod_{i=j}^t (1 - \gamma_i \lambda_i).$$

In particular, if the two sequences $(\gamma_t)_{t \in \mathbb{N}}$ and $(\underline{\alpha}_t)_{t \in \mathbb{N}}$ are such that, for all $t \geq j_0$, $\gamma_t \lambda_t := c \bar{t}^{-1}$ for some $c, t_0 > 0$, then (B) yields

$$\|\Pi_j^t\| \leq \left(\frac{j+t_0}{\bar{t}+1} \right)^c.$$

3.3. Application to online learning and Proof of Theorem A. The online learning sequence $(f_t)_{t \in \mathbb{N}_0}$ defined in (6), with assumptions (I)-(II), can be interpreted as a sequential stochastic approximation algorithm $(w_t)_{t \in \mathbb{N}_0}$ in (8), taking values in the Hilbert space $\mathcal{W} := \mathcal{H}_K$, with

$$\begin{aligned} A((x, y)) &:= \langle \cdot, K_x \rangle_K K_x, \quad b((x, y)) := y K_x \\ A_t &:= A + \lambda_t I, \quad b_t := b, \end{aligned}$$

so that

$$\begin{aligned} \bar{A} &= L_K, \quad \bar{b} = L_K f_\rho, \quad \bar{w} = f_\rho, \\ \bar{A}_t &= L_K + \lambda_t I, \quad \bar{w}_t = f_{\lambda_t}, \\ r_t &= f_t - f_{\lambda_t}, \quad \Delta_t = f_{\lambda_t} - f_{\lambda_{t-1}}. \end{aligned}$$

Let us emphasize that the operator A is only defined from \mathcal{H}_K to \mathcal{H}_K here (we would not be able to define $f(x)$ for $f \in \mathcal{L}_{\rho, \mathcal{X}}^2$). The properties mentioned below will only hold on \mathcal{H}_K in general, and in particular the norms of operators $\|\cdot\|$ are assumed to be $\|\cdot\|_{\mathcal{H}_K \rightarrow \mathcal{H}_K}$, although operators defined on $\mathcal{L}_{\rho, \mathcal{X}}^2$ and commuting with $L_K^{1/2}$ (which is an isometry between $\mathcal{L}_{\rho, \mathcal{X}}^2$ and H_K) have the same norm in either spaces.

Note that $A(z)$ is positive for all $z = (x, y) \in \mathcal{X}$ (which implies $\bar{A} = L_K$ positive as well), since

$$\langle A((x, y))(f), f \rangle = \langle f(x) K_x, f \rangle = f(x)^2 \geq 0$$

for all $f \in \mathcal{H}_K$.

Also, for all $f \in \mathcal{H}_K$, $\|A f\| = |\langle K_x, f \rangle| \|K_x\| \leq \|K_x\|^2 \|f\|$, so that

$$\|A\| \leq \kappa^2, \quad \|\bar{A}\| \leq \mathbb{E}(\|A\|) \leq \kappa^2.$$

Hence

$$(16) \quad \|A_t\| \leq \bar{\alpha}_t := \lambda_t + \kappa^2, \quad \|A_t^{-1}\|^{-1} \geq \underline{\alpha}_t := \lambda_t.$$

Let us prove Theorem A; assume its conditions hold. Then the *Generalized Finiteness Condition* of Section 3.2 is satisfied. Now $f_\rho \in \mathcal{H}_K$ implies $\|f_\lambda - f_\rho\|_K \rightarrow 0$ when $\lambda \rightarrow 0$. Therefore the conclusion follows from the convergence of $\mathbb{E}[\|w_t - \bar{w}_t\|^2] = \mathbb{E}[\|f_t - f_{\lambda_t}\|^2]$ to 0 in Theorem 3.5, the condition of uniform boundedness of $\mathbb{E}\|A_t \bar{w}_t - b_t\|^2$ being shown in Lemma 5.5 (B).

For convenience, we will use, in Sections 5 and 6, the notation

$$L_t := A(z_t) = \langle \cdot, K_{x_t} \rangle_K K_{x_t}$$

We will assume that

$$(17) \quad \gamma_t = \frac{a}{\bar{t}^\theta}, \quad \lambda_t = \frac{b}{\bar{t}^{1-\theta}}, \quad \text{for some } \theta \in [0, 1], t_0 > 0, a \in (0, t_0^\theta), b \in (0, t_0^{1-\theta}),$$

and then study the \mathcal{H}_K or $\mathcal{L}_{\rho, \mathcal{X}}^2$ -norm of the error $f_t - f_\rho$, based using a reverse martingale (resp. martingale) decomposition in Section 5 (resp. Section 6). We will then optimize the upper bounds in θ , a and b by using some prior information on the regularity of f_ρ .

Finally observe that Lemma 3.8 implies, using (16), that for all $j, t \in \mathbb{N}$, $t \geq j$,

$$(18) \quad \|I - \gamma_t A_t\| \leq 1 - \gamma_t \lambda_t = 1 - \frac{ab}{t}, \quad \|\Pi_j^t\| \leq \left(\frac{j+t_0}{t+1}\right)^{ab}$$

if $t_0^\theta \geq a(\kappa^2 + b)$ (and, therefore, $\gamma_t \bar{\alpha}_t = \gamma_t \lambda_t + \gamma_t \kappa^2 \leq abt_0^{-1} + a\kappa^2 t_0^{-\theta} \leq 1$).

Similarly, for all $j, t \in \mathbb{N}$, $t \geq j$,

$$(19) \quad \|I - \gamma_t \bar{A}_t\| \leq 1 - \frac{ab}{t}, \quad \|\bar{\Pi}_j^t\| \leq \left(\frac{j+t_0}{t+1}\right)^{ab}$$

if $t_0^\theta \geq a(\kappa^2 + b)$; the norm in (19) can be $\|\cdot\|_{\mathcal{H}_K \rightarrow \mathcal{H}_K}$ as well as $\|\cdot\|_{\mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{L}_{\rho, \mathcal{X}}^2}$.

4. ESTIMATES OF DRIFT ON THE REGULARIZATION PATH

This section is devoted to estimates on the drift $\|f_\lambda - f_\mu\|$, $\lambda, \mu > 0$, along the regularization path $\lambda \rightarrow f_\lambda$, in \mathcal{H}_K -norm or $\mathcal{L}_{\rho, \mathcal{X}}^2$ -norm, assuming that $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r > 0$. These estimates enable us to upper bound on the one hand the approximation error $\|f_\lambda - f_\rho\|$ (when specialized to $\mu = 0$), and on the other hand the *drift error* in the martingale and reversed martingale decompositions.

Note that the estimate $\|f_\lambda - f_\mu\|_K = O(|\lambda - \mu|)$ in the case $r = 1$ is not improved by increasing r . This is related to a phenomenon usually referred to as the *saturation* problem in regularizations [Engl, Hanke, and Neubauer 1996].

Theorem 4.1. *Let $\lambda > \mu \geq 0$. Assume that $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r \geq -1$.*

(A) *If $r \in [-1, 1] \setminus \{0\}$, then*

$$\|f_\lambda - f_\mu\|_\rho \leq |\lambda^r - \mu^r| \frac{\|L_K^{-r} f_\rho\|_\rho}{r};$$

(B) *If $r \geq 1$, then for any $1 \leq s \leq r$,*

$$\|f_\lambda - f_\mu\|_\rho \leq \kappa^{2(s-1)} |\lambda - \mu| \|L_K^{-s} f_\rho\|_\rho;$$

(C) *If $r \geq 1/2$, then*

$$\|f_\lambda - f_\mu\|_K \leq \frac{|\lambda - \mu|}{\lambda} \|f_\rho\|_K;$$

(D) *If $r \in [-1/2, 3/2] \setminus \{1/2\}$, then*

$$\|f_\lambda - f_\mu\|_K \leq |\lambda^{r-1/2} - \mu^{r-1/2}| \frac{\|L_K^{-r} f_\rho\|_\rho}{|r - \frac{1}{2}|};$$

(E) *If $r \geq 3/2$, then for any $3/2 \leq s \leq r$,*

$$\|f_\lambda - f_\mu\|_K \leq \kappa^{2(s-3/2)} |\lambda - \mu| \|L_K^{-s} f_\rho\|_\rho.$$

Proof. Fix $\lambda > \mu$, assume $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r \in [-1, 1]$ and let $\|\cdot\| := \|\cdot\|_{\mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{L}_{\rho, \mathcal{X}}^2}$. We first prove that, for all $u \geq -1$, if we let

$$J_{u, \lambda, \mu} := (\mu - \lambda)(L_K + \lambda I)^{-1}(L_K + \mu I)^{-1}L_K^{1+u}$$

then, for all $t \in [-1, 1] \setminus \{0\}$, $u \geq t$,

$$(20) \quad \|J_{u, \lambda, \mu}\| \leq \kappa^{2(u-t)} |\lambda^t - \mu^t|/|t|.$$

This will be useful, since

$$(21) \quad f_\lambda - f_\mu = (\mu - \lambda)(L_K + \lambda I)^{-1}(L_K + \mu I)^{-1}L_K f_\rho = J_{r, \lambda, \mu} L_K^{-r} f_\rho,$$

using that

$$(L_K + \lambda I)f_\lambda = L_K f_\rho, \quad (L_K + \mu I)f_\mu = L_K f_\rho.$$

Let us prove (20): using $\|L_K^{u-t}\| = \|L_K\|^{u-t} \leq \kappa^{2(u-t)}$ by (3), and $\max(t, 0) + \min(0, t) = t$,

$$\begin{aligned} \|J_{u, \lambda, \mu}\| &\leq |\lambda - \mu| \|(L_K + \lambda I)^{\max(t, 0)-1}(L_K + \mu I)^{\min(t, 0)} L_K^{-(t+1)} L_K^{1+u}\| \\ &\leq |\lambda - \mu| \lambda^{-1} \lambda^{\max(t, 0)} \mu^{\min(t, 0)} \|L_K^{u-t}\| \leq \kappa^{2(u-t)} |\lambda - \mu| \lambda^{-1} \max(\lambda^t, \mu^t) \\ &= \kappa^{2(u-t)} \Lambda(\mu) |\lambda^t - \mu^t|, \end{aligned}$$

where

$$\Lambda(\mu) := \begin{cases} \frac{1-\mu/\lambda}{1-(\mu/\lambda)^t} & \text{if } t > 0 \\ \frac{1-\mu/\lambda}{1-(\lambda/\mu)^t} & \text{if } t < 0 \end{cases}$$

Now

$$\Lambda(\mu) \leq \frac{1}{|t|}.$$

Indeed, if $t > 0$, then this is a consequence of $x \leq (1 - (1 - x)^t)/t$ applied to $x := 1 - \mu/\lambda$, using that $x \mapsto (1 - (1 - x)^t)/t$ (defined on $(-\infty, 1]$) is convex and thus remains above the tangent line at 0. Similarly, we use $x \leq (1 - (1 - x)^{-t})/(-t)$ if $t < 0$.

Now (20)-(21) implies (A) with $u := r$ and $t := r$, and (B) with $u := s$ and $t := 1$, since $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ implies $L_K^{-s} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for any $s \leq r$. Similarly, (D) (resp. (E)) follows from

$$L_K^{-1/2} (f_\lambda - f_\mu) = J_{r-1/2, \lambda, \mu} L_K^{-r} f_\rho,$$

and (20) applied to $u := r - 1/2$ and $t := u$ (resp. $u := s - 1/2$ and $t := 1$).

Let us now prove (C): if $r \geq 1/2$, then $f_\rho \in \mathcal{H}_K$, and the first part of equality (21) implies

$$\|f_\lambda - f_\mu\|_K \leq |\mu - \lambda| \|(L_K + \lambda I)^{-1}\| \|(L_K + \mu I)^{-1} L_K\| \|L_K^{-1/2} f_\rho\|_\rho \leq \frac{|\mu - \lambda|}{\lambda} \|f_\rho\|_K.$$

□

5. UPPER BOUNDS FOR CONVERGENCE IN \mathcal{H}_K

Throughout this section, we assume that $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{H}}^2$ for some $r \in (1/2, 3/2]$, which implies $f_\rho \in \mathcal{H}_K$ with additional regularity, and assume that the sequences $(\gamma_t)_{t \in \mathbb{N}}$ and $(\lambda_t)_{t \in \mathbb{N}}$ are chosen in (17).

Our goal is to provide a probabilistic upper bound for

$$\|f_t - f_\rho\|_K,$$

in order to prove Theorem B. We start with the triangle inequality

$$\|f_t - f_\rho\|_K \leq \|f_t - f_{\lambda_t}\|_K + \|f_{\lambda_t} - f_\rho\|_K,$$

and apply the reversed martingale decomposition of $(f_t)_{t \in \mathbb{N}}$ developed in Section 3, Theorem 3.1:

$$(22) \quad r_t = \Pi_1^t r_0 - \sum_{j=1}^t \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j) - \sum_{j=1}^t \Pi_j^t \Delta_j.$$

We make use of the corresponding notation of Section 3, in particular Section 3.3, so that

$$A_j \bar{w}_j - b_j = (L_t + \lambda_t I) f_{\lambda_t} - y_t K_{x_t},$$

and

$$\Pi_j^t(x_j, \dots, x_t) = \begin{cases} \prod_{i=j}^t (I - \gamma_i (L_i + \lambda_i I)) & \text{if } j \leq t; \\ I & \text{otherwise.} \end{cases}$$

Now

$$\|f_t - f_\rho\|_K \leq \mathcal{E}_{init}(t) + \mathcal{E}_{samp}(t) + \mathcal{E}_{drift}(t) + \mathcal{E}_{approx}(t),$$

where we define the errors as follows:

- (A) *Initial Error*: $\mathcal{E}_{init}(t) = \|\Pi_1^t r_0\|_K$ comes from the initial choice f_0 ;
- (B) *Approximation Error*: $\mathcal{E}_{approx}(t) = \|f_{\lambda_t} - f_\rho\|_K$, measures the distance between the regression function and the regularization path at time t ;
- (C) *Drift Error*: $\mathcal{E}_{drift}(t) = \|\sum_{j=1}^t \Pi_j^t \Delta_j\|_K$ comes from the drift along the regularization path $t \mapsto f_{\lambda_t}$;
- (D) *Sample Error*: $\mathcal{E}_{samp}(t) = \|\sum_{j=1}^t \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j)\|_K$, where $\xi_j = \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j)$ is a reversed martingale difference sequence, reflecting the random fluctuation caused by sampling.

In the remainder of this section, we are going to provide upper bounds for each of the four errors, which, roughly speaking when $ab = 1$, are

$$\begin{aligned} \mathcal{E}_{init}(t) &= O(t^{-1}), \\ \mathcal{E}_{approx}(t) &= O(t^{-(r-1/2)(1-\theta)}), \\ \mathcal{E}_{drift}(t) &= O(t^{-(r-1/2)(1-\theta)}), \\ \mathcal{E}_{samp}(t) &= O(t^{\frac{1}{2}-\theta}). \end{aligned}$$

It is not surprising that the approximation error and drift error have the same rate, as both of them come from the estimates on drifts in Theorem 4.1. This suggests our explanation that the bias = $\mathcal{E}_{approx}(t) + \mathcal{E}_{drift}(t)$ and the variance = $\mathcal{E}_{samp}(t)$. Theorem B then follows from these bounds by setting $\theta = 2r/(2r+1)$.

5.1. Initial Error.

Theorem 5.1 (Initial Error). *Let $t_0^\theta \geq a(\kappa^2 + b)$. Then for all $t \in \mathbb{N}$,*

$$\mathcal{E}_{init}(t) \leq B_3 \bar{t}^{-ab},$$

where $B_3 = (t_0 + 1)^{ab} \|r_0\|$.

Proof.

$$\mathcal{E}_{init}(t) \leq \|\Pi_1^t\| \|r_0\| \leq \left(\frac{t_0 + 1}{\bar{t} + 1} \right)^{ab} \|r_0\| \leq \left(\frac{t_0 + 1}{\bar{t}} \right)^{ab} \|r_0\|$$

where the second last step uses Lemma 3.8 (B) with $j = 1$. \square

5.2. Approximation Error. The approximation error is derived from Theorem 4.1(D) by setting $\lambda = \lambda_t$ and $\mu = 0$.

Theorem 5.2 (Approximation Error). *For $r \in (1/2, 3/2]$ and $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$,*

$$\|f_{\lambda_t} - f_\rho\|_K \leq B_1 b^{r-1/2} \bar{t}^{-(r-1/2)(1-\theta)},$$

where $B_1 = (r - 1/2)^{-1} \|L_K^{-r} f_\rho\|_\rho$.

5.3. Drift Error.

Theorem 5.3 (Drift Error). *Let $t_0^\theta \geq [a(\kappa^2 + b) \vee 1]$. Then for $r \in (1/2, 3/2]$ and $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$,*

$$\mathcal{E}_{drift}(t) \leq \begin{cases} B_2 b^{r-1/2} \bar{t}^{-(r-1/2)(1-\theta)}, & \text{if } ab > (r - 1/2)(1 - \theta); \\ B_2 b^{r-1/2} \bar{t}^{-ab}, & \text{if } ab < (r - 1/2)(1 - \theta), \end{cases}$$

where $B_2 = \frac{4(1 - \theta)}{|ab - (r - 1/2)(1 - \theta)|} \|L_K^{-r} f_\rho\|_\rho$.

Proof. We are going to provide an upper bound of

$$\mathcal{E}_{drift}(t) = \left\| \sum_{j=1}^t \Pi_j^t \Delta_j \right\|_K.$$

First, Lemma 3.8 implies, using (16), that for all $j, t \in \mathbb{N}, t \geq j$,

$$(23) \quad \|\Pi_j^t\| \leq \left(\frac{j + t_0}{\bar{t} + 1} \right)^{ab},$$

if $t_0^\theta \geq a(\kappa^2 + b)$ (and, therefore, $\gamma_t \bar{\alpha}_t = \gamma_t \lambda_t + \gamma_t \kappa^2 \leq ab t_0^{-1} + a \kappa^2 t_0^{-\theta} \leq 1$).

Second, by Theorem 4.1(D),

$$(24) \quad \begin{aligned} \|\Delta_t\|_K &= \|f_{\lambda_t} - f_{\lambda_{t-1}}\|_K \leq |\lambda_t^{r-1/2} - \lambda_{t-1}^{r-1/2}| \frac{\|L_K^{-r} f_\rho\|_\rho}{r - \frac{1}{2}} \\ &\leq b^{r-1/2} (1 - \theta) (\bar{t} - 1)^{-(r-1/2)(1-\theta)-1} \|L_K^{-r} f_\rho\|_\rho, \end{aligned}$$

where we use

$$\begin{aligned} |\lambda_t^{r-1/2} - \lambda_{t-1}^{r-1/2}| &= b^{r-1/2} \left| \bar{t}^{-(r-1/2)(1-\theta)} - (\bar{t}-1)^{-(r-1/2)(1-\theta)} \right| \\ &\leq b^{r-1/2} (r-1/2)(1-\theta) (\bar{t}-1)^{-(r-1/2)(1-\theta)-1}, \end{aligned}$$

due to the Mean Value Theorem with $h(x) = x^{-(r-1/2)(1-\theta)}$ and $h'(x) = -(r-1/2)(1-\theta)x^{-(r-1/2)(1-\theta)-1}$, such that

$$|h\bar{t} - h(\bar{t}-1)| = |h'(\eta)| \leq |h'(\bar{t}-1)|, \quad \text{for some } \eta \in (\bar{t}-1, \bar{t}).$$

Now combining (23) and (24) gives

$$\begin{aligned} \mathcal{E}_{drift}(t) &= \left\| \sum_{j=1}^t \Pi_j^t \Delta_j \right\|_K \leq b^{r-1/2} (1-\theta) \|L_K^{-r} f_\rho\|_\rho \cdot \sum_{j=1}^t \left(\frac{j+t_0}{\bar{t}+1} \right)^{ab} (j+t_0-1)^{-(r-1/2)(1-\theta)-1} \\ &\leq \frac{4b^{r-1/2} (1-\theta) \|L_K^{-r} f_\rho\|_\rho}{(\bar{t}+1)^{ab}} \sum_{j=1}^t (j+t_0)^{ab-1-(r-1/2)(1-\theta)}. \end{aligned}$$

It suffices to bound

$$\sum_{j=1}^t (j+t_0)^{ab-1-(r-1/2)(1-\theta)} \leq \int_0^{t+1} (x+t_0)^{ab-1-(r-1/2)(1-\theta)} dx =: I_t$$

Now, if $ab > (r-1/2)(1-\theta)$, then

$$I_t \leq \frac{(\bar{t}+1)^{ab-(r-1/2)(1-\theta)}}{ab - (r-1/2)(1-\theta)};$$

whereas $ab < (r-1/2)(1-\theta)$ implies

$$I_t \leq \frac{t_0^{ab-(r-1/2)(1-\theta)}}{(r-1/2)(1-\theta) - ab} \leq \frac{1}{|ab - (r-1/2)(1-\theta)|}, \quad t_0 \geq 1.$$

□

5.4. Sample Error.

Theorem 5.4 (Sample Error). *Assume that $t_0^\theta \geq [a(\kappa^2 + b) \vee b \vee 1]$ and $ab \neq \theta - 1/2$ or $(3\theta - 1)/2$. Then, with probability at least $1 - \delta$ ($\delta \in (0, 1)$),*

$$\mathcal{E}_{samp}(t) \leq B_4 ab^{-1/2} \bar{t}^{-[ab \wedge \frac{3\theta-1}{2}]} + B_5 a \bar{t}^{-[ab \wedge (\theta-1/2)]}$$

where $B_4 = \frac{2(\kappa+1)^2 M_\rho}{3} \log \frac{2}{\delta}$ and $B_5 = \frac{8\kappa M_\rho}{\sqrt{|ab - (\theta-1/2)|}} \log \frac{2}{\delta}$.

The proof of Theorem 5.4 requires some auxiliary estimates.

Lemma 5.5. *Let $A_t \bar{w}_t - b_t = (f_{\lambda_t}(x_t) - y_t) K_{x_t} + \lambda_t f_{\lambda_t}$.*

(A) $\|A_t \bar{w}_t - b_t\|_K \leq (\kappa+1)^2 M_\rho / \sqrt{\lambda_t}$, if $t_0^{1-\theta} \geq b$;

(B) $\mathbb{E}[\|A_t \bar{w}_t - b_t\|_K^2] \leq 4\kappa^2 M_\rho^2$.

Proof. (A) Using $\|f_\lambda\|_K \leq M_\rho/\sqrt{\lambda}$ in Lemma B.1(A),

$$\|A_t \bar{w}_t - b_t\| \leq \|f_{\lambda_t}(x_t)K_{x_t}\|_K + |y_t| \|K_{x_t}\|_K + \lambda_t \|f_{\lambda_t}\|_K \leq M_\rho \kappa^2 / \sqrt{\lambda_t} + M_\rho \kappa + M_\rho \sqrt{\lambda_t}$$

since $\|f_{\lambda_t}(x_t)K_{x_t}\|_K = |\langle f_{\lambda_t}, K_{x_t} \rangle| \|K_{x_t}\|_K \leq \|f_{\lambda_t}\|_K \|K_{x_t}\|_K^2 \leq M_\rho \kappa^2 / \sqrt{\lambda_t}$. Now,

$$M_\rho \kappa^2 / \sqrt{\lambda_t} + M_\rho \kappa + M_\rho \sqrt{\lambda_t} \leq (\kappa^2 + \kappa + 1) M_\rho / \sqrt{\lambda_t} \leq (\kappa + 1)^2 M_\rho / \sqrt{\lambda_t}$$

where the second last inequality is due to $t_0^{1-\theta} \geq b \Rightarrow \lambda_t \leq 1$.

(B) Using $\lambda_t f_\lambda = L_K f_\rho - L_K f_\lambda$ we obtain

$$(f_{\lambda_t}(x_t) - y_t) K_{x_t} + \lambda_t f_{\lambda_t} = (L_t - L_K) f_{\lambda_t} + L_K f_\rho - y_t K_{x_t}.$$

$$\begin{aligned} \mathbb{E}[\|A_t \bar{w}_t - b_t\|^2] &= \mathbb{E}\|(L_t - L_K) f_{\lambda_t} + L_K f_\rho - y_t K_{x_t}\|_K^2 \\ &\leq 2\mathbb{E}[\|(L_t - L_K) f_{\lambda_t}\|_K^2 + \|L_K f_\rho - y_t K_{x_t}\|_K^2] \\ &\leq 2\mathbb{E}[\|L_t f_{\lambda_t}\|_K^2 + \|y_t K_{x_t}\|_K^2] \leq 2\kappa^2(\|f_{\lambda_t}\|_\rho^2 + M_\rho^2) = 4\kappa^2 M_\rho^2 \end{aligned}$$

since $\mathbb{E}[L_t] = L_K$, $\mathbb{E}[y_t K_{x_t}] = L_K f_\rho$ and $\|f_\lambda\|_\rho \leq M_\rho$ by Lemma B.1(B). \square

Now we are ready to give the proof of the sample error bounds, Theorem 5.4.

Proof of Theorem 5.4. We are going to bound

$$\mathcal{E}_{\text{samp}}(t) = \left\| \sum_{j=1}^t \xi_j \right\|_K$$

where $\xi_j = \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j)$ is a reversed martingale difference sequence. To apply the Pinelis-Bernstein inequality in Proposition A.3, we need bounds on $\|\xi_j\|_K$ and $\mathbb{E}_{j+1} \|\xi_j\|_K^2$ where $\mathbb{E}_{j+1}[\cdot]$ is the expectation conditional on examples after time j .

Notice that for $t_0^\theta \geq a(\kappa^2 + b)$ and $j \geq 1$, using $1 + x \leq e^x$ for all $x \in \mathbb{R}$,

$$\|\gamma_j \Pi_{j+1}^t\| \leq \frac{a}{(j+t_0)^\theta} \left(\frac{j+t_0+1}{\bar{t}+1} \right)^{ab} \leq \frac{a(j+t_0)^{ab-\theta}}{(\bar{t}+1)^{ab}} (1+t_0^{-1})^{ab} \leq \frac{ea(j+t_0)^{ab-\theta}}{(\bar{t}+1)^{ab}},$$

where e is the Euler constant.

Now Lemma 5.5 (B) implies

$$\mathbb{E}[\|A_t \bar{w}_t - b_t\|_K^2] \leq 4\kappa^2 M_\rho^2.$$

Hence

$$\mathbb{E}_{j+1} \|\xi_j\|^2 \leq \frac{4e^2(a\kappa M_\rho)^2 (j+t_0)^{2ab-2\theta}}{(\bar{t}+1)^{2ab}},$$

so that, if $t_0 \geq 2$,

$$(25) \quad \sum_{j=1}^t \mathbb{E}_{j+1} \|\xi_j\|^2 \leq \begin{cases} \frac{2e^2(a\kappa M_\rho)^2}{ab - (\theta - 1/2)} (\bar{t}+1)^{-2(\theta-1/2)}, & \text{if } ab > \theta - 1/2 \\ \frac{2e^2(a\kappa M_\rho)^2}{(\theta - 1/2) - ab} (\bar{t}+1)^{-2ab}, & \text{if } ab < \theta - 1/2 \end{cases}$$

On the other hand, if $t_0^{1-\theta} \geq b$, Lemma 5.5 (A) implies

$$\|A_j \bar{w}_j - b_j\|_K \leq (\kappa + 1)^2 M_\rho / \sqrt{\lambda_j},$$

whence

$$\begin{aligned}
 \|\xi_j\|_K &\leq \frac{ea(\kappa+1)^2 M_\rho}{\sqrt{b}} \cdot \frac{(j+t_0)^{ab-(3\theta-1)/2}}{(\bar{t}+1)^{ab}} \\
 (26) \quad &\leq \begin{cases} \frac{ea(\kappa+1)^2 M_\rho}{\sqrt{b}} \bar{t}^{-(3\theta-1)/2}, & \text{if } ab \geq (3\theta-1)/2 \\ \frac{ea(\kappa+1)^2 M_\rho}{\sqrt{b}} \bar{t}^{-ab}, & \text{if } ab \leq (3\theta-1)/2 \end{cases}
 \end{aligned}$$

The final bound is obtained by Pinelis-Bernstein inequality in Proposition A.3 with upper bounds (25) and (26). \square

5.5. Proof of Theorem B. We choose $\theta = 2r/(2r+1)$, $a \geq 1$, $b \leq 1$ such that $ab = 1$, and assume $t_0^\theta \geq a\kappa^2 + 1$. Using Theorems 5.2, 5.3, 5.1, and 5.4,

$$\begin{aligned}
 \|f_t - f_\rho\|_K &\leq \mathcal{E}_{init}(t) + \mathcal{E}_{approx}(t) + \mathcal{E}_{drift}(t) + \mathcal{E}_{samp}(t) \\
 &\leq B_3 \bar{t}^{-ab} + [(B_1 + B_2)a^{1/2-r} + (B_4 t_0^{-\theta} \sqrt{a} + B_5)a] \bar{t}^{-(2r-1)/(4r+2)}
 \end{aligned}$$

Note that, by Lemma 5.5(A) with $f_0 = 0$,

$$B_3 = (t_0 + 1) \|r_0\| = (t_0 + 1) \|f_{\lambda_0}\| \leq C_0 := 2t_0 \frac{M_\rho}{\sqrt{\lambda_0}} = 2t_0^{\frac{4r+3}{4r+2}} M_\rho$$

On the other hand,

$$C_1 := B_1 + B_2 = \left(\frac{2}{2r-1} + \frac{8}{2r+3} \right) \|L_K^{-r} f_\rho\|_\rho = \frac{20r-2}{(2r-1)(2r+3)} \|L_K^{-r} f_\rho\|_\rho$$

and, using $\sqrt{a}t_0^{-\theta} \leq \kappa^{-1}$ and $t_0 \geq 1$,

$$B_4 t_0^{-\theta} \sqrt{a} + B_5 \leq \frac{2(\kappa+1)^2 M_\rho}{3\kappa} + \frac{8\kappa M_\rho}{\sqrt{3/4}} \leq C_2 := \frac{20(\kappa+1)^2 M_\rho}{\kappa},$$

which concludes the proof of Theorem B.

6. UPPER BOUNDS FOR CONVERGENCE IN $\mathcal{L}_{\rho, \mathcal{X}}^2$

Throughout this section, we assume that $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r \in [1/2, 3/2]$, which implies $f_\rho \in \mathcal{H}_K$ with additional regularity, and assume the sequences $(\gamma_t)_{t \in \mathbb{N}}$ and $(\lambda_t)_{t \in \mathbb{N}}$ are chosen in (17). Note that the case $r = 1/2$ is included here, whereas it was not in Section 5 and Theorem B.

Our goal is to provide a probabilistic upper bound of

$$\|f_t - f_\rho\|_\rho,$$

in order to prove Theorem C. As in Section 5, we start with the triangle inequality

$$\|f_t - f_\rho\|_\rho \leq \|f_t - f_{\lambda_t}\|_\rho + \|f_{\lambda_t} - f_\rho\|_\rho,$$

but apply here the martingale decomposition of $(f_t)_{t \in \mathbb{N}_0}$ in $\mathcal{L}_{\rho, \mathcal{X}}^2$ developed in Theorem 3.3 instead:

$$r_t = \bar{\Pi}_1^t r_0 + \sum_{j=1}^t \gamma_j \bar{\Pi}_{j+1}^t \chi_j - \sum_{j=1}^t \bar{\Pi}_j^t \Delta_j.$$

We make use of the corresponding notation of Section 3, in particular 3.3, so that

$$\chi_t = (L_K - L_t)f_{t-1} + (y_t K_{x_t} - L_K f_\rho),$$

and

$$(27) \quad \bar{\Pi}_j^t = \begin{cases} \prod_{i=j}^t (I - \gamma_i(L_K + \lambda_i I)) & \text{if } j \leq t; \\ I & \text{otherwise.} \end{cases}$$

The martingale decomposition enables us to make use of the isometry $L_K^{1/2} : \mathcal{L}_{\rho, \mathcal{X}}^2 \rightarrow \mathcal{H}_K$, in the sense that one can benefit from the spectral decomposition of $L_K^{1/2} \bar{\Pi}_j^t$ to get a tighter estimate. This was not possible with the reversed martingale decomposition, since $L_K^{1/2} \Pi_j^t$ does not have an obvious spectral decomposition.

Note however that χ_t depends on f_{t-1} , so that we need preliminary estimates of $\|\chi_t\|_\rho$, provided in Appendix B.

As in Section 5, we introduce the following definitions for convenience.

[Definitions of Errors]

- (A) *Initial Error*: $\mathcal{E}_{init}(t) = \|\bar{\Pi}_1^t r_0\|_\rho$, which reflects the propagation error by the initial choice f_0 ;
- (B) *Approximation Error*: $\mathcal{E}_{approx}(t) = \|f_{\lambda_t} - f_\rho\|_\rho$, which measures the distance between the regression function and the regularization path at time t ;
- (C) *Drift Error*: $\mathcal{E}_{drift}(t) = \|\sum_{j=1}^t \bar{\Pi}_j^t \Delta_j\|_\rho$, which measures the error caused by drifts from $f_{\lambda_{j-1}}$ to f_{λ_j} along the regularization path;
- (D) *Sample Error*: $\mathcal{E}_{samp}(t) = \|\sum_{j=1}^t \gamma_j \bar{\Pi}_{j+1}^t \chi_j\|_\rho$, where χ_j is a martingale difference sequence, reflecting the random fluctuation caused by sampling.

Our aim is to bound

$$\|f_t - f_\rho\|_\rho \leq \mathcal{E}_{init}(t) + \mathcal{E}_{samp}(t) + \mathcal{E}_{drift}(t) + \mathcal{E}_{approx}(t).$$

In the remainder of this section, we are going to provide upper bounds for each of the four errors, which, roughly speaking when $ab = 1$, are

$$\begin{aligned} \mathcal{E}_{init}(t) &= O(t^{-1}) \\ \mathcal{E}_{approx}(t) &= O(t^{-r(1-\theta)}) \\ \mathcal{E}_{drift}(t) &= O(t^{-r(1-\theta)}) \\ \mathcal{E}_{samp}(t) &= O(t^{-\theta/2}) \end{aligned}$$

This suggests our explanation that the bias $= \mathcal{E}_{approx}(t) + \mathcal{E}_{drift}(t) = O(t^{-r(1-\theta)})$ and the variance $= \mathcal{E}_{samp}(t) = O(t^{-\theta/2})$ similar to the batch learning setting. Theorem C then follows from these bounds by setting $\theta = 2r/(2r+1)$.

6.1. Initial Error.

Theorem 6.1 (Initial Error). *Let $t_0^\theta \geq a(\kappa^2 + b)$. Then for all $t \in \mathbb{N}$,*

$$\mathcal{E}_{init}(t) \leq B_6 \bar{t}^{-ab},$$

where $B_6 = M_\rho(t_0 + 1)^{ab}$.

Proof. Lemma 3.8(B) with $j = 1$ and (16) imply that, if $t_0^\theta \geq a(\kappa^2 + b)$,

$$\mathcal{E}_{init}(t) \leq \|\bar{\Pi}_1^t\| \|r_0\| \leq \left(\frac{t_0 + 1}{t + 1}\right)^{ab} \|r_0\|.$$

For $f_0 = 0$, using Lemma B.1(B), $\|r_0\|_\rho = \|f_{\lambda_0}\|_\rho \leq M_\rho$. \square

6.2. Approximation Error.

Theorem 6.2 (Approximation Error). *For $r \in (0, 1]$ and $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$,*

$$\mathcal{E}_{approx}(t) \leq B_7 b^r \bar{t}^{-r(1-\theta)},$$

where $B_7 = r^{-1} \|L_K^{-r} f_\rho\|_\rho$.

Proof. Follows from Theorem 4.1(A) with $\lambda = \lambda_t$ and $\mu = 0$. \square

6.3. Drift Error.

Theorem 6.3 (Drift Error). *Assume $t_0^\theta \geq [a(\kappa^2 + b) \vee 1]$. Then, if $r \in (0, 1]$ and $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$,*

$$\mathcal{E}_{drift}(t) \leq \begin{cases} B_8 b^r \bar{t}^{-r(1-\theta)}, & \text{if } ab > r(1-\theta); \\ B_8 b^r \bar{t}^{-ab}, & \text{if } ab < r(1-\theta), \end{cases}$$

where $B_8 = \frac{4(1-\theta)}{|ab - r(1-\theta)|} \|L_K^{-r} f_\rho\|_\rho$.

Proof. Similar to the proof of Theorem 5.3, replacing $r - 1/2$ with r . \square

6.4. Sample Error. In this section we assume $b = a^{-1}$ for simplicity; this is necessary for the bounds in Appendix B, in particular Corollary B.7, and is enough to provide the optimal bounds we need (see discussion after statement of Theorem A).

Theorem 6.4 (Sample Error). *Assume that $L_K^{-r} f_\rho \in \mathcal{L}_{\rho, \mathcal{X}}^2$ for some $r \in [1/2, 1]$, $\theta \in [1/2, 2/3]$, $ab = 1$, $a \geq 4$ and $t_0^\theta \geq 2 + 8\kappa^2 a$. Then, for all $t \in \mathbb{N}$, with probability at least $1 - \delta$ ($\delta \in (0, 1)$)*

$$\mathcal{E}_{samp}(t) \leq \sqrt{a} B_9 \frac{\log(2/\delta)}{\bar{t}^{\theta/2}} + \left(a^{3/2} B_{10} \sqrt{\log \bar{t}} + a^{5/2} B_{11}\right) \frac{(\log(2/\delta))^2}{\bar{t}^{(3\theta-1)/2}},$$

where

$$B_9 := 10\kappa M_\rho, \quad B_{10} := 63\kappa^2 M_\rho, \quad B_{11} := 50\kappa^2 M_\rho t_0^{1/2-\theta}.$$

Proof. Fix $t \in \mathbb{N}$, $\delta \in [0, 1]$, and let

$$(28) \quad A_{t, \delta} := \kappa M_\rho a \left[12at_0^{1/2-\theta} + 15\sqrt{\log \bar{t}} \right] \log \frac{2}{\delta}.$$

For all $j \in \mathbb{N}$, let us define the martingale difference sequence

$$X_j := \gamma_j \bar{\Pi}_{j+1}^t \chi_j \mathbf{1}_{\{\|h_{j-1}\|_K (j+t_0)^{\theta-1/2} \leq A_{t, \delta}\}},$$

where we make use of the notation of Appendix B. Recall that Corollary B.7 implies, with probability at least $1 - \delta$, all the indicator function events for $1 \leq j < t$ hold, which will be assumed in the computation below.

Recall that

$$\chi_j = (\bar{A}_j - A_j)w_{j-1} + b_j - \bar{b}_j = (L_K - L_j)f_{j-1} + y_j K_{x_j} - L_K f_\rho$$

where $L_j := \langle \cdot, K_{x_j} \rangle K_{x_j}$.

Using Lemma B.3 and the decomposition $f_j = f_\rho + g_j + h_j$ in Appendix B, we deduce that, for all $1 \leq j < t$, if $\|h_{j-1}\|_K (j + t_0)^{\theta-1/2} \leq A_{t,\delta}$,

$$\begin{aligned} \mathbb{E}_{j-1} \|\chi_j\|_K^2 &= \mathbb{E}_{j-1} \|y_j K_{x_j} - L_j f_{j-1} - (L_K f_\rho - L_K f_{j-1})\|_K^2 \leq \mathbb{E}_{j-1} \|y_j K_{x_j} - L_j f_{j-1}\|_K^2 \\ &\leq \mathbb{E}_{j-1} \|(y_j K_{x_j} - L_j f_\rho) - L_j g_{j-1} - L_j h_{j-1}\|_K^2 \\ &\leq 3\kappa^2 [\mathbb{E}_{j-1} |y_j - f_\rho(x_j)|^2 + \mathbb{E}_{j-1} |g_{j-1}(x_j)|^2 + \mathbb{E}_{j-1} |h_{j-1}(x_j)|^2] \\ &\leq 3\kappa^2 [4M_\rho^2 + \|g_{j-1}\|_\rho^2 + \|h_{j-1}\|_\rho^2] \leq 3\kappa^2 [5M_\rho^2 + \kappa^2 (j + t_0)^{1-2\theta} A_{t,\delta}^2] =: A'_{j,t,\delta} \end{aligned}$$

Now, using the isometry $L_K^{1/2} : \mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{H}_K$,

$$\begin{aligned} \sum_{j=1}^t \mathbb{E}_{j-1} \|X_j\|_\rho^2 &= \sum_{j=1}^t \mathbb{E}_{j-1} \|L_K^{1/2} X_j\|_K^2 = \sum_{j=1}^t \gamma_j^2 \mathbb{E}_{j-1} \|L_K^{1/2} \bar{\Pi}_{j+1}^t \chi_j\|_K^2 \\ &\leq \sum_{j=1}^t (\gamma_j^2 \|\bar{\Pi}_{j+1}^t L_K \bar{\Pi}_{j+1}^t\|) \mathbb{E}_{j-1} \|\chi_j\|_K^2 \leq \sum_{j=1}^t \gamma_j^2 A'_{j,t,\delta} \|\bar{\Pi}_{j+1}^t L_K \bar{\Pi}_{j+1}^t\| \end{aligned}$$

In order to estimate $\sum_{j=1}^t \gamma_j^2 A'_{j,t,\delta} \|\bar{\Pi}_{j+1}^t L_K \bar{\Pi}_{j+1}^t\|$, recall that $(\mu_\alpha, \phi_\alpha)_{\alpha \in \mathbb{N}}$ is an orthonormal eigen-system of $L_K : \mathcal{L}_{\rho_{\mathcal{X}}}^2 \rightarrow \mathcal{L}_{\rho_{\mathcal{X}}}^2$. Let $a_i = \gamma_i \lambda_i + \gamma_i \mu_\alpha$ for simplicity; then

$$\begin{aligned} \sum_{j=1}^t \gamma_j^2 A'_{j,t,\delta} \|\bar{\Pi}_{j+1}^t L_K \bar{\Pi}_{j+1}^t\| &\leq \sup_{\mu_\alpha} \sum_{j=1}^t \gamma_j^2 A'_{j,t,\delta} \mu_\alpha \prod_{i=j+1}^t (1 - a_i)^2 \\ &= \sup_{\mu_\alpha} \sum_{j=1}^t \left[\gamma_j A'_{j,t,\delta} \prod_{i=j+1}^t (1 - a_i) \right] \cdot \left[\gamma_j \mu_\alpha \prod_{i=j+1}^t (1 - a_i) \right] \\ &\leq \sup_{\mu_\alpha} \left\{ \left[\sup_j \gamma_j A'_{j,t,\delta} \prod_{i=j+1}^t (1 - a_i) \right] \cdot \left[\sum_{j=1}^t \gamma_j \mu_\alpha \prod_{i=j+1}^t (1 - a_i) \right] \right\} \end{aligned}$$

where for large enough t_0 ,

$$\begin{aligned} \sup_j \gamma_j A'_{j,t,\delta} \prod_{i=j+1}^t (1 - a_i) &\leq \sup_j \gamma_j A'_{j,t,\delta} \prod_{i=j+1}^t (1 - \gamma_i \lambda_i) \\ &\leq 3a\kappa^2 \sup_j \frac{j + t_0}{\bar{t}} \cdot \left(\frac{5M_\rho^2}{(j + t_0)^\theta} + \frac{\kappa^2 A_{t,\delta}^2}{(j + t_0)^{3\theta-1}} \right) \\ &\leq 3a\kappa^2 \left(\frac{5M_\rho^2}{\bar{t}^\theta} + \frac{\kappa^2 A_{t,\delta}^2}{\bar{t}^{3\theta-1}} \right), \end{aligned}$$

and

$$\sum_{j=1}^t \gamma_j \mu_\alpha \prod_{i=j+1}^t (1 - a_i) \leq \sum_{j=1}^t (1 - (1 - \gamma_j \mu_\alpha)) \prod_{i=j+1}^t (1 - \gamma_i \mu_\alpha) = 1 - \prod_{i=1}^t (1 - \gamma_i \mu_\alpha) \leq 1.$$

These two upper bounds give

$$(29) \quad \sum_{j=1}^t \mathbb{E}_{j-1} \|X_j\|_\rho^2 \leq \frac{3a\kappa^2}{\bar{t}^\theta} \left(5M_\rho^2 + \frac{\kappa^2 A_{t,\delta}^2}{\bar{t}^{2\theta-1}} \right) =: \sigma_t^2.$$

Moreover, again if $\|h_{j-1}\|_K (j+t_0)^{\theta-1/2} \leq A_{t,\delta}$, then, using Lemma B.3 (B) and Corollary B.7, we deduce

$$\begin{aligned} \|y_j K_{x_j} - L_j f_{j-1}\|_K &= \|y_j K_{x_j} - L_j(f_\rho + g_{j-1} + h_{j-1})\|_K \\ &\leq \kappa M_\rho + \frac{\kappa^2 M_\rho}{\sqrt{\lambda_j}} + \kappa^2 A_{t,\delta} (j+t_0)^{1/2-\theta} =: C_{j,t,\delta}, \end{aligned}$$

which implies

$$\|\chi_j\|_K = \|y_j K_{x_j} - L_j f_{j-1} - \mathbb{E}_j[y_j K_{x_j} - L_j f_{j-1}]\|_K \leq 2C_{j,t,\delta}.$$

Therefore

$$\begin{aligned} \|X_j\|_\rho &\leq \gamma_j \|L_K^{1/2} \bar{\Pi}_{j+1}^t \chi_j\|_K \leq 2\gamma_j C_{j,t,\delta} \|\bar{\Pi}_{j+1}^t L_K \bar{\Pi}_{j+1}^t\|_K^{1/2} \\ &\leq 2\kappa \sup_j \gamma_j C_{j,t,\delta} \prod_{i=j+1}^t (1 - \gamma_i \lambda_i), \quad \|L_K\| \leq \kappa^2 \\ &\leq 2a\kappa^2 \sup_j \frac{j+t_0}{\bar{t}} \cdot \left(\frac{M_\rho}{(j+t_0)^\theta} + \frac{\kappa M_\rho \sqrt{a}}{(j+t_0)^{(3\theta-1)/2}} + \frac{\kappa A_{t,\delta}}{(j+t_0)^{2\theta-1/2}} \right) \\ &\leq 2a\kappa^2 \left(\frac{M_\rho}{\bar{t}^\theta} + \frac{\kappa M_\rho \sqrt{a}}{\bar{t}^{(3\theta-1)/2}} + \frac{\kappa A_{t,\delta}}{\bar{t}^{2\theta-1/2}} \right) \\ &\leq \frac{2\sqrt{a}\kappa}{\bar{t}^{\theta/2}} \left(M_\rho + \kappa \frac{\kappa M_\rho a + A_{t,\delta}}{\bar{t}^{\theta-1/2}} \right) =: M, \end{aligned}$$

where we use $t_0 \geq \kappa^2$ twice in the last inequality.

Combining M and σ_t from (29), we obtain

$$\begin{aligned} 2 \left(\frac{M}{3} + \sigma_t \right) &= \frac{2\sqrt{a}\kappa}{\bar{t}^{\theta/2}} \left[\left(\sqrt{15} + \frac{2}{3} \right) M_\rho + \kappa \frac{(\sqrt{3} + 1/3) A_{t,\delta} + \kappa M_\rho a / 3}{\bar{t}^{\theta-1/2}} \right] \\ &\leq \frac{\sqrt{a} B_9}{\bar{t}^{\theta/2}} + \left(a^{3/2} B_{10} \sqrt{\log \bar{t}} + a^{5/2} B_{11} \right) \frac{\log(2/\delta)}{\bar{t}^{(3\theta-1)/2}}, \end{aligned}$$

where we use that

$$\begin{aligned} B_9 &= 10\kappa M_\rho \geq 2\kappa(\sqrt{15} + 2/3)M_\rho, \\ B_{10} &= 63\kappa^2 M_\rho \geq \kappa^2 M_\rho [30(\sqrt{3} + 1/3) + 2/(3\log 2)], \\ B_{11} &= 50\kappa^2 M_\rho t_0^{1/2-\theta} \geq 24\kappa^2 M_\rho t_0^{1/2-\theta} (\sqrt{3} + 1/3). \end{aligned}$$

□

6.5. Proof of Theorem C. We choose $\theta = 2r/(2r+1)$, $a \geq 1$, $b \leq 1$ such that $ab = 1$, and assume $t_0^\theta \geq a\kappa^2 + 1$. Using Theorems, 6.2, 6.3, 6.1, and 6.4

$$\begin{aligned} \|f_t - f_\rho\|_\rho &\leq \mathcal{E}_{init}(t) + \mathcal{E}_{approx}(t) + \mathcal{E}_{drift}(t) + \mathcal{E}_{samp}(t) \\ &\leq \frac{B_6}{t} + \left((B_7 + B_8)a^{-r} + \sqrt{a}B_9 \log \frac{2}{\delta} \right) \left(\frac{1}{t} \right)^{\frac{r}{2r+1}} + \left(a^{3/2}B_{10}\sqrt{\log t} + a^{5/2}B_{11} \right) \frac{(\log(2/\delta))^2}{t^{\frac{6r-1}{4r+2}}} \end{aligned}$$

This enables us to conclude, with $D_0 := 2M_\rho t_0 \geq B_6 = M_\rho(t_0 + 1)$,

$$D_1 := B_7 + B_8 = \frac{5r+1}{r(1+r)} \|L_K^{-r} f_\rho\|_\rho,$$

$D_2 := B_9$, $D_3 := B_{10}$, and $D_4 := B_{11}$.

APPENDIX A: A PROBABILISTIC INEQUALITY

The following result is quoted from [Theorem 3.4 in Pinelis 1994].

Lemma A.1 (Pinelis-Bennett). *Let ξ_i be a martingale difference sequence in a Hilbert space. Suppose that almost surely $\|\xi_i\| \leq M$ and $\sum_{i=1}^t \mathbb{E}_{i-1} \|\xi_i\|^2 \leq \sigma_t^2$. Then*

$$\mathbf{Prob} \left\{ \sup_{1 \leq k \leq t} \left\| \sum_{i=1}^k \xi_i \right\| \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\sigma_t^2}{M^2} g \left(\frac{M\epsilon}{\sigma_t^2} \right) \right\},$$

where $g(x) = (1+x) \log(1+x) - x$ for $x > 0$.

Using the lower bound $g(x) \geq \frac{x^2}{2(1+x/3)}$, one may obtain the following generalized Bernstein's inequality.

Corollary A.2 (Pinelis-Bernstein). *Let ξ_i be a martingale difference sequence in a Hilbert space. Suppose that almost surely $\|\xi_i\| \leq M$ and $\sum_{i=1}^t \mathbb{E}_{i-1} \|\xi_i\|^2 \leq \sigma_t^2$. Then*

$$(A-1) \quad \mathbf{Prob} \left\{ \sup_{1 \leq k \leq t} \left\| \sum_{i=1}^k \xi_i \right\| \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\epsilon^2}{2(\sigma_t^2 + M\epsilon/3)} \right\}.$$

The following result will be used as a basic probabilistic inequality to derive various bounds.

Proposition A.3. *Let ξ_i be a martingale difference sequence in a Hilbert space. Suppose that almost surely $\|\xi_i\| \leq M$ and $\sum_{i=1}^t \mathbb{E}_{i-1} \|\xi_i\|^2 \leq \sigma_t^2$. Then the following holds with probability at least $1 - \delta$ ($\delta \in (0, 1)$),*

$$\sup_{1 \leq k \leq t} \left\| \sum_{i=1}^k \xi_i \right\| \leq 2 \left(\frac{M}{3} + \sigma_t \right) \log \frac{2}{\delta}.$$

Proof. Taking the right hand side of (A-1) to be δ , then we arrive at the following quadratic equation for ϵ ,

$$\epsilon^2 - \frac{2M}{3}\epsilon \log \frac{2}{\delta} - 2\sigma_t^2 \log \frac{2}{\delta} = 0.$$

Note that $\epsilon > 0$, then

$$\begin{aligned}\epsilon &= \frac{1}{2} \left\{ \frac{2M}{3} \log \frac{2}{\delta} + \sqrt{\frac{4M^2}{9} \log^2 \frac{2}{\delta} + 8\sigma_t^2 \log \frac{2}{\delta}} \right\} \\ &= \frac{M}{3} \log \frac{2}{\delta} + \sqrt{\left(\frac{M}{3}\right)^2 \log^2 \frac{2}{\delta} + 2\sigma_t^2 \log \frac{2}{\delta}} \\ &\leq \frac{2M}{3} \log \frac{2}{\delta} + \sqrt{2\sigma_t^2 \log \frac{2}{\delta}},\end{aligned}$$

where the second last step is due to $\sqrt{a^2 + b^2} \leq a + b$ ($a, b > 0$) with

$$a = \frac{M}{3} \log \frac{2}{\delta}, \quad \text{and} \quad b = \sqrt{2\sigma_t^2 \log \frac{2}{\delta}}.$$

We complete the proof by relaxing $\sqrt{2\sigma_t^2 \log 2/\delta} \leq 2\sigma_t \log 2/\delta$ since $2 \log 2/\delta > 1$ for $\delta \in (0, 1)$. \square

APPENDIX B: PRELIMINARY UPPER BOUNDS

Appendix B is devoted to the proof of preliminary upper bounds on the online learning sequence $(f_t)_{t \in \mathbb{N}}$ defined in (6), and on the regularization path $\lambda \mapsto f_\lambda$. We make use of the notation of Section 3, in particular Section 3.3. For simplicity we assume $f_0 := 0$; note that another choice would correspond to adding $\Pi_1^t f_0$ to f_t at time t . We assume that the sequences $(\gamma_t)_{t \in \mathbb{N}}$ and $(\lambda_t)_{t \in \mathbb{N}}$ are chosen as in (17).

Firstly, Lemmas B.1 and B.2 provide deterministic upper bounds. Then the rest of the Appendix aims at obtaining probabilistic bounds of $(f_t)_{t \in \mathbb{N}}$, based on a decomposition of $f_t - f_\rho$ into two parts in (B-2): g_t is purely deterministic and is upper bounded in $\mathcal{L}_{\rho, \mathcal{X}}^2$ -norm in Lemma B.3, and h_t is studied in detail in Lemmas B.4 and following. Lemma B.7 yields logarithmic estimates with large probability.

Lemma B.1. *For any $\lambda > 0$,*

- (A) $\|f_\lambda\|_K \leq M_\rho / \sqrt{\lambda}$;
- (B) $\|f_\lambda\|_\rho \leq M_\rho$.

Proof. (A) By definition,

$$f_\lambda = \arg \min_{f \in \mathcal{H}_K} \|f - f_\rho\|_\rho^2 + \lambda \|f\|_K^2.$$

The term we minimize on the right-hand side takes the value $\|f_\rho\|_\rho^2$ at $f = 0$, so that

$$(B-1) \quad \|f_\lambda - f_\rho\|_\rho^2 + \lambda \|f_\lambda\|_K^2 \leq \|f_\rho\|_\rho^2 \leq M_\rho^2,$$

which yields the result.

(B) Using (5),

$$\|f_\lambda\|_\rho = \|(L_K + \lambda I)^{-1} L_K f_\rho\|_\rho \leq \|(L_K + \lambda I)^{-1} L_K\| \|f_\rho\|_\rho \leq \|f_\rho\|_\rho \leq M_\rho.$$

\square

Lemma B.2. Assume $t_0^\theta \geq a(\kappa^2 + b)$. Then, for all $t \in \mathbb{N}$,

$$\|f_t\|_K \leq \frac{\kappa M_\rho}{\lambda_t}.$$

Proof. Recall that $f_t = (I - \gamma_t A_t)f_{t-1} + \gamma_t y_t K_{x_t}$. Now assume $t_0^\theta \geq \kappa^2 + 1$: using (18),

$$\|f_t\|_K \leq \|1 - \gamma_t A_t\| \|f_{t-1}\|_K + \gamma_t \|y_t K_{x_t}\|_K \leq (1 - \gamma_t \lambda_t) \|f_{t-1}\|_K + \gamma_t \kappa M_\rho.$$

By induction on t , we deduce

$$\|f_t\|_K \leq \kappa M_\rho \sum_{j=1}^t \gamma_j \prod_{i=j+1}^t (1 - \gamma_i \lambda_i) \leq \max_{1 \leq j \leq t} \left(\frac{1}{\lambda_j} \right) \sum_{j=1}^t \gamma_j \lambda_j \prod_{i=j+1}^t (1 - \gamma_i \lambda_i) \leq \frac{1}{\lambda_t},$$

since

$$\sum_{j=1}^t \gamma_j \lambda_j \prod_{i=j+1}^t (1 - \gamma_i \lambda_i) = 1 - \prod_{i=1}^t (1 - \gamma_i \lambda_i).$$

□

In the rest of Appendix, we prove probabilistic bounds of $(f_t)_{t \in \mathbb{N}_0}$. First observe that the definition of the online learning sequence (6) can be rewritten as

$$f_t - f_\rho = [I - \gamma_t(L_t + \lambda_t I)](f_{t-1} - f_\rho) + \gamma_t(y_t K_{x_t} - L_t f_\rho) - \gamma_t \lambda_t f_\rho.$$

Let us now define the following $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ -adapted processes $(g_t)_{t \in \mathbb{N}_0}$ and $(h_t)_{t \in \mathbb{N}_0}$ recursively by

$$g_0 := -f_\rho, \quad h_0 := 0,$$

and

$$\begin{aligned} g_t &:= [I - \gamma_t(L_t + \lambda_t I)]g_{t-1} - \gamma_t \lambda_t f_\rho, \\ h_t &:= [I - \gamma_t(L_t + \lambda_t I)]h_{t-1} + \gamma_t(y_t K_{x_t} - L_t f_\rho) + \gamma_t(L_K - L_t)g_{t-1}. \end{aligned}$$

We can easily prove by induction that

$$(B-2) \quad f_t - f_\rho = g_t + h_t,$$

using $f_0 = 0$.

Lemma B.3. Assume $t_0^\theta \geq a(\kappa^2 + b)$. Then, for all $t \in \mathbb{N}_0$,

- (A) $\|g_t\|_\rho \leq M_\rho$;
- (B) $\|g_t + f_\rho\|_K \leq 3M_\rho/\sqrt{\lambda_t}$.

Proof. We prove (A) by induction: $\|g_0\|_\rho = \|f_\rho\|_\rho \leq M_\rho$ and, for all $t \in \mathbb{N}$, if we assume $\|g_{t-1}\|_\rho \leq M_\rho$ then, using (19),

$$\|g_t\|_\rho \leq \|I - \gamma_t(L_K + \lambda_t I)\| \|g_{t-1}\|_\rho + \gamma_t \lambda_t \|f_\rho\|_\rho \leq (1 - \gamma_t \lambda_t) \|g_{t-1}\|_\rho + \gamma_t \lambda_t M_\rho \leq M_\rho.$$

To prove (B), observe that, for all $t \in \mathbb{N}$,

$$\begin{aligned} g_t + f_\rho &= [I - \gamma_t(L_K + \lambda_t I)]g_{t-1} + (1 - \gamma_t \lambda_t) f_\rho = [I - \gamma_t(L_K + \lambda_t I)](g_{t-1} + f_\rho) + \gamma_t L_K f_\rho \\ &= [I - \gamma_t(L_K + \lambda_t I)](g_{t-1} + f_\rho) + \gamma_t(L_K + \lambda_t I) f_{\lambda_t}, \end{aligned}$$

so that

$$g_t + f_\rho - f_{\lambda_t} = [I - \gamma_t(L_K + \lambda_t I)](g_{t-1} + f_\rho - f_{\lambda_t}).$$

Let, for all $t \in \mathbb{N}$,

$$w_t := g_t + f_\rho - f_{\lambda_t}.$$

Then it is easy to show by induction that

$$w_t = \bar{\Pi}_1^t w_0 + \sum_{k=1}^t \bar{\Pi}_k^t (f_{\lambda_k} - f_{\lambda_{k-1}})$$

which implies, using Theorem 4.1 (D) with $r = 0$, and Lemma B.1 (A) ($w_0 = -f_{\lambda_0}$) that

$$\|w_t\| \leq 2M_\rho / \sqrt{\lambda_t}.$$

This enables us to conclude, using again Lemma B.1 (A). \square

For all $t \in \mathbb{N}_0$ and $M \in \mathbb{R}_+ \cup \{\infty\}$, let

$$\begin{aligned} \bar{L}_t &:= \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}} L_t, & \tilde{L}_t &:= \mathbf{1}_{\{|h_{t-1}(x_t)| > M\}} L_t, \\ \bar{L}_K &:= \mathbb{E}_{t-1}[\bar{L}_t], & \tilde{L}_K &:= \mathbb{E}_{t-1}[\tilde{L}_t]. \end{aligned}$$

Note that $L_t = \bar{L}_t + \tilde{L}_t$ and $L_K = \bar{L}_K + \tilde{L}_K$.

For all $t \in \mathbb{N}$, let

$$(B-3) \quad \bar{h}_t := [I - \gamma_t(\bar{L}_t + \lambda_t I)]h_{t-1} + \gamma_t(y_t K_{x_t} - L_t f_\rho) + \gamma_t(L_K - L_t)g_{t-1} = h_t + \gamma_t \tilde{L}_t h_{t-1}$$

$$(B-4) \quad k_t := \bar{h}_t - (1 - \gamma_t \lambda_t)h_{t-1} = \gamma_t[-\bar{L}_t h_{t-1} + (y_t K_{x_t} - L_t f_\rho) + (L_K - L_t)g_{t-1}]$$

$$(B-5) \quad = \gamma_t[-\bar{L}_t h_{t-1} + y_t K_{x_t} + L_K g_{t-1} - L_t(f_\rho + g_{t-1})].$$

In Lemma B.4 we upper bound $\|\bar{h}_t\|_K^2$ in conditional expectation; note that the result still holds when $M = \infty$. We threshold h_t into \bar{h}_t in order to limit its conditional variance, which will be necessary in order to obtain logarithmic estimates with large probability in Lemma B.7, using on the other hand Lemma B.6 showing that, if M is large enough, $\|h_t\|_K \leq \|\bar{h}_t\|_K$.

Lemma B.4. *Assume $t_0^\theta \geq 2a(b + \kappa^2)$. For all $M \in \mathbb{R}_+ \cup \{\infty\}$ and $t \in \mathbb{N}$,*

$$\mathbb{E}_{t-1}[\|\bar{h}_t\|_K^2] \leq (1 - \gamma_t \lambda_t)^2 \|h_{t-1}\|_K^2 + 3\kappa^2 M_\rho^2 \gamma_t^2.$$

In particular, assume moreover that $\theta \geq 1/2$, $\epsilon := ab - (\theta - 1/2) > 0$ and $t_0 \geq \max(2ab, 2\epsilon, \epsilon + (2\theta - 1)/\epsilon)$, and let $A := a\kappa M_\rho \sqrt{3/\epsilon}$. Then $\|h_{t-1}\|_K \geq A\bar{t}^{1/2-\theta}$ implies

$$\bar{t}^{\theta-1/2} \mathbb{E}_{t-1}[\|\bar{h}_t\|_K] \leq (\bar{t} - 1)^{\theta-1/2} \|h_{t-1}\|_K.$$

Proof. For all $t \in \mathbb{N}$, let

$$\zeta_t := (\bar{L}_K - \bar{L}_t)h_{t-1} + (L_K - L_t)g_{t-1} + (y_t K_{x_t} - L_t f_\rho),$$

so that

$$\bar{h}_t = [I - \gamma_t(\bar{L}_K + \lambda_t I)]h_{t-1} + \gamma_t \zeta_t.$$

Using $\mathbb{E}_{t-1}[\zeta_t] = 0$, we deduce that

$$(B-6) \quad \mathbb{E}_{t-1}[\|\bar{h}_t\|_K^2] = \|[I - \gamma_t(\bar{L}_K + \lambda_t I)]h_{t-1}\|_K^2 + \gamma_t^2 \mathbb{E}_{t-1}[\|\zeta_t\|_K^2].$$

Let us now upper bound the two summands in the right-hand side of equality (B-6). First,

$$\begin{aligned} & \| [I - \gamma_t (\bar{L}_K + \lambda_t I)] h_{t-1} \|_K^2 \\ &= (1 - \gamma_t \lambda_t)^2 \| h_{t-1} \|_K^2 - 2\gamma_t (1 - \gamma_t \lambda_t) \mathbb{E}_{t-1} [|h_{t-1}(x_t)|^2 \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}}] \\ &\quad + \gamma_t^2 \mathbb{E}_{t-1} [\bar{L}_t h_{t-1}] \|_K^2, \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}_{t-1} [\bar{L}_t h_{t-1}] \|_K^2 &\leq (\mathbb{E}_{t-1} [|h_{t-1}(x_t)| \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}} \|K_{x_t}\|_K])^2 \\ &\leq \kappa^2 \mathbb{E}_{t-1} [|h_{t-1}(x_t)|^2 \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}}], \end{aligned}$$

using conditional Jensen's inequality.

Second, using that $\mathbb{E}[y_t K_{x_t} - L_t f_\rho \mid \sigma(\mathcal{F}_{t-1}, x_t)] = 0$,

$$\begin{aligned} \mathbb{E}_{t-1} [\|\zeta_t\|_K^2] &= \mathbb{E}_{t-1} [\|(\bar{L}_K - \bar{L}_t) h_{t-1} + (L_K - L_t) g_{t-1}\|_K^2 + \|y_t K_{x_t} - L_t f_\rho\|_K^2] \\ &\leq \mathbb{E}_{t-1} [\|\bar{L}_t h_{t-1} + L_t g_{t-1}\|_K^2 + \|y_t K_{x_t}\|_K^2] \\ &\leq \kappa^2 (2\mathbb{E}_{t-1} [|h_{t-1}(x_t)|^2 \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}}] + 2\|g_{t-1}\|_\rho^2 + M_\rho^2) \\ &\leq \kappa^2 (2\mathbb{E}_{t-1} [|h_{t-1}(x_t)|^2 \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}}] + 3M_\rho^2), \end{aligned}$$

where we use Lemma B.3 (A) in the last inequality.

In summary, we obtain that

$$\begin{aligned} \mathbb{E}_{t-1} [\|\bar{h}_t\|_K^2] &\leq (1 - \gamma_t \lambda_t)^2 \|h_{t-1}\|_K^2 - \gamma_t (2 - 2\gamma_t \lambda_t - 3\gamma_t \kappa^2) \mathbb{E}_{t-1} [|h_{t-1}(x_t)|^2 \mathbf{1}_{\{|h_{t-1}(x_t)| \leq M\}}] \\ &\quad + 3\kappa^2 M_\rho^2 \gamma_t^2. \end{aligned}$$

Now, the assumption $t_0^\theta \geq 2a(b + \kappa^2)$ implies $2 - 2\gamma_t \lambda_t - 3\gamma_t \kappa^2 \geq 0$ for all $t \in \mathbb{N}$, which completes the proof of the first statement.

Let us now prove the second statement:

$$\begin{aligned} \Delta_t &:= \mathbb{E}_{t-1} \left[\left(1 - \frac{1}{\bar{t}}\right)^{1-2\theta} \|\bar{h}_t\|^2 - \|h_{t-1}\|^2 \right] \\ (B-7) \quad &\leq \left(1 - \frac{1}{\bar{t}}\right)^{1-2\theta} \left(1 - \frac{ab}{\bar{t}}\right)^2 \|h_{t-1}\|^2 - \|h_{t-1}\|^2 + 3\kappa^2 M_\rho^2 a^2 \bar{t}^{-2\theta}. \end{aligned}$$

Now, since $t_0 \geq \max(2ab, 2\epsilon, \epsilon + (2\theta - 1)/\epsilon)$ and $\theta \geq 1/2$, we have

$$\log \left[\left(1 - \frac{1}{\bar{t}}\right)^{1-2\theta} \left(1 - \frac{ab}{\bar{t}}\right)^2 \left(1 - \frac{\epsilon}{\bar{t}}\right)^{-1} \right] \leq -\frac{\epsilon}{\bar{t}} + \frac{2\theta - 1 + \epsilon^2}{\bar{t}^2} \leq 0.$$

using $\log(1 - x) \leq -x$ for all $x \in [0, 1]$ and $\log(1 - x) \geq -x - x^2$ for all $x \in [0, 1/2]$.

Therefore (B-7) implies

$$\Delta_t \leq -\frac{\epsilon}{\bar{t}} \|h_{t-1}\|^2 + 3\kappa^2 M_\rho^2 a^2 \bar{t}^{-2\theta} \leq 0.$$

The conclusion follows by conditional Jensen's inequality. \square

Lemma B.5. *Assume $t_0^\theta \geq a(\kappa^2 + b)$ and $t_0^{1-\theta} \geq b(2 + MM_\rho^{-1})$; then*

$$\|k_t\|_K \leq 2\kappa M_\rho ab^{-1} \bar{t}^{1-2\theta} \text{ and } \mathbb{E}_{t-1} [\|k_t\|_K^2] \leq 9\gamma_t^2 M_\rho^2 \kappa^2.$$

Proof. By definition (B-5), using $\|K_{x_t}\|_K \leq \kappa$, $\|L_K f_\rho\|_K \leq \kappa \|L_K^{1/2} f_\rho\|_K = \kappa \|f_\rho\|_\rho \leq \kappa M_\rho$ and Lemma B.3 (A)-(B), we deduce

$$\|k_t\|_K \leq \gamma_t [\kappa(M + 2M_\rho) + \|L_t(f_\rho + g_{t-1})\|_K] \leq \kappa \gamma_t \left(M + 2M_\rho + \frac{M_\rho}{\lambda_t} \right) \leq \frac{2\kappa \gamma_t M_\rho}{\lambda_t} = \frac{2\kappa M_\rho ab^{-1}}{\bar{t}^{2\theta-1}},$$

where we use $t_0^{1-\theta} \geq b(2 + MM_\rho^{-1})$ in the last inequality.

Now, using (B-4), we obtain

$$\begin{aligned} \mathbb{E}_{t-1}[\|k_t\|_K^2] &\leq 3\gamma_t^2 [\mathbb{E}_{t-1}[\|\bar{L}_t h_{t-1}\|_K^2] + \mathbb{E}_{t-1}[\|y_t K_{x_t}\|_K^2] + \mathbb{E}_{t-1}[\|L_t g_{t-1}\|_K^2]] \\ &\leq 3\gamma_t^2 [2M_\rho^2 \kappa^2 + \|g_{t-1}\|_\rho \kappa^2] \leq 9\gamma_t^2 M_\rho^2 \kappa^2. \end{aligned}$$

□

Lemma B.6. For all $t \in \mathbb{N}$, assume $M \geq M_t := 4\kappa M_\rho ab^{-1}\bar{t}^{1-2\theta}$, $t_0^\theta \geq 2a(\kappa^2 + b)$ and $t_0^{1-\theta} \geq b(2 + MM_\rho^{-1})$; then

$$\|h_t\|_K \leq \|\bar{h}_t\|_K.$$

Proof. Assume $h_{t-1}(x_t) \geq M_t$ for instance; the other case is similar. By definition,

$$h_t = \bar{h}_t - \gamma_t h_{t-1}(x_t) K_{x_t}$$

so that

$$\|h_t\|_K^2 = \|\bar{h}_t\|_K^2 - 2\gamma_t h_{t-1}(x_t) \bar{h}_t(x_t) + \gamma_t^2 (h_{t-1}(x_t))^2 K(x_t, x_t) \leq \|\bar{h}_t\|_K^2$$

if $\bar{h}_t(x_t) \geq \kappa^2 \gamma_t h_{t-1}(x_t)/2$.

But, using Lemma B.5,

$$\bar{h}_t(x_t) = (1 - \gamma_t \lambda_t) h_{t-1}(x_t) + k_t(x_t) \geq (1 - \gamma_t \lambda_t) h_{t-1}(x_t) - 2\kappa M_\rho ab^{-1}\bar{t}^{1-2\theta} \geq \kappa^2 \gamma_t h_{t-1}(x_t)/2$$

if

$$2\kappa M_\rho ab^{-1}\bar{t}^{1-2\theta} \leq h_{t-1}(x_t)/2 \leq h_{t-1}(x_t)(1 - \gamma_t \lambda_t - \kappa^2 \gamma_t/2),$$

since the assumption $t_0^\theta \geq 2a(\kappa^2 + b)$ implies $1 - \gamma_t \lambda_t - \kappa^2 \gamma_t/2 \geq 1/2$. □

The following logarithmic upper bound holds under the assumptions $ab - (\theta - 1/2) > 0$, $\theta \in [1/2, 1]$ and t_0 sufficiently large, but we assume $b = a^{-1}$ in its statement, for notational reasons.

Corollary B.7. Assume $\theta \in [1/2, 1]$, $b = a^{-1}$, $t_0^\theta \geq 2 + 8\kappa^2 a$ and $t_0^{1-\theta} \geq 4b$. Then, with probability at least $1 - \delta$,

$$\sup_{0 \leq k \leq t} \|h_k\|_K (k + t_0 + 1)^{\theta-1/2} \leq \kappa M_\rho a \left[12a t_0^{1/2-\theta} + 15\sqrt{\log \bar{t}} \right] \log \frac{2}{\delta}.$$

Proof. Let us first check that the assumptions of Lemmas B.4, B.5 and B.6 are satisfied, and apply these lemmas: $t_0^\theta \geq 3 + 8\kappa^2 a \geq 2a(\kappa^2 + b)$. Now $\epsilon = ab - (\theta - 1/2) \in [1/2, 1]$, and the hypothesis $t_0 \geq \max(2ab, 2\epsilon, \epsilon + (2\theta - 1)/\epsilon)$ is satisfied as long $t_0 \geq 3$, which is assumed here. We choose $M = M_t = 4\kappa M_\rho ab^{-1}\bar{t}^{1-2\theta}$; now $t_0^\theta \geq 8\kappa a$ and $t_0^{1-\theta} \geq 4b$ imply $t_0^{1-\theta} \geq b(2 + 4\kappa ab^{-1}t_0^{1-2\theta}) \geq b(2 + M_t M_\rho^{-1})$.

For all $i \in \mathbb{N}$, if $\|h_{i-1}\|_K \geq A(i+t_0)^{1/2-\theta}$, $A := a\kappa M_\rho \sqrt{3/\epsilon}$, then

$$\begin{aligned} \|h_i\|_K &\leq \|\bar{h}_i\|_K, && \text{(Lemma B.6)} \\ &\leq \|\bar{h}_i\|_K + \mathbb{E}_{i-1}(\|\bar{h}_i\|_K) - \mathbb{E}_{i-1}(\|\bar{h}_i\|_K) \\ &\leq \left(1 - \frac{1}{i+t_0}\right)^{\theta-1/2} \|h_{i-1}\|_K + \epsilon_i, && \text{(Lemma B.4)} \end{aligned}$$

where

$$\epsilon_i := \|\bar{h}_i\|_K - \mathbb{E}_{i-1}(\|\bar{h}_i\|_K)$$

satisfies

$$\|\epsilon_i\|_K \leq 4\kappa M_\rho ab^{-1}(i+t_0)^{1-2\theta} \text{ and } \mathbb{E}_{i-1}[\|\epsilon_i\|^2] \leq 9\gamma_i^2 M_\rho^2 \kappa^2.$$

Let, for all $i \in \mathbb{N}$,

$$\eta_i := \sum_{k=1}^i \epsilon_k (k+t_0)^{\theta-1/2} \mathbf{1}_{\{\|h_{k-1}\|_K \geq A(k+t_0)^{1/2-\theta}\}}.$$

Fix $t \in \mathbb{N}$. For all $0 \leq i < t$, $\|\eta_{i+1} - \eta_i\| \leq 4\kappa M_\rho a^2 t_0^{1/2-\theta}$, and

$$\sum_{k=1}^t \mathbb{E}_{k-1} \|\eta_k\|^2 \leq 9\kappa^2 M_\rho^2 a^2 \sum_{k=1}^t (k+t_0)^{-1} \leq 9\kappa^2 M_\rho^2 a^2 \log(1+t/t_0).$$

Let

$$\Delta := \left\{ \sup_{1 \leq i \leq t} \|\eta_i\| \leq 2\kappa M_\rho a \left[\frac{4at_0^{1/2-\theta}}{3} + 3\sqrt{\log\left(1 + \frac{t}{t_0}\right)} \right] \log \frac{2}{\delta} \right\}.$$

By Proposition A.3, $P(\Delta) \geq 1 - \delta$.

Now assume Δ holds. Let, for all $k \in \mathbb{N}$,

$$x_k := \|h_k\|_K (k+t_0+1)^{\theta-1/2}.$$

For all $k \leq t$, let

$$m := \max\{j \leq k : \|h_j\|_K < A(j+t_0+1)^{1/2-\theta}\}.$$

If $m < k$, then

$$\begin{aligned} x_{m+1} &\leq [A(m+t_0+1)^{1/2-\theta} + 2\kappa M_\rho a^2 (m+t_0+1)^{1-2\theta}] (m+t_0+2)^{\theta-1/2} \\ &\leq \frac{\sqrt{5}}{2} [A + 2\kappa M_\rho a^2 t_0^{1/2-\theta}] \leq \frac{\sqrt{5}}{2} [\sqrt{6}a\kappa M_\rho + 2\kappa M_\rho a^2 t_0^{1/2-\theta}]; \end{aligned}$$

the second inequality comes from $[(m+t_0+2)/(m+t_0+1)]^{\theta-1/2} \leq \sqrt{5/4}$, since $t_0 \geq 3$.

On the other hand it is easy to prove by induction that, for all $k \leq t$,

$$x_k \leq x_{m+1} + \eta_k - \eta_{m+1}$$

and, therefore,

$$\begin{aligned}
x_k &\leq \kappa M_\rho a \left[\left(\sqrt{5} + \frac{16}{3} \right) a t_0^{1/2-\theta} + \frac{\sqrt{30}}{2} + 12 \sqrt{\log \left(1 + \frac{t}{t_0} \right)} \right] \log \frac{2}{\delta} \\
&\leq 12 \kappa M_\rho a \left[a t_0^{1/2-\theta} + \frac{1}{4} + \sqrt{\log \left(1 + \frac{k}{t_0} \right)} \right] \log \frac{2}{\delta} \\
&\leq \kappa M_\rho a \left[12 a t_0^{1/2-\theta} + 15 \sqrt{\log t} \right] \log \frac{2}{\delta},
\end{aligned}$$

using in the last inequality that, for all $t \geq 1$ and $t_0 \geq 2$,

$$\frac{1}{4} + \sqrt{\log(1 + t/t_0)} \leq \frac{1}{4} + \sqrt{\log(t + t_0)} \leq \frac{5}{4} \sqrt{\log(t + t_0)}.$$

□

APPENDIX C: PROOF OF RESULTS OF SECTION 3.2

Proof of Lemma 3.8. Assume $t \geq j_0$. The spectral Theorem for compact operators implies that there is an orthonormal basis of \mathcal{W} consisting of eigenvectors of A_t , so that, if $(\alpha_{t,k})_{k \in \mathbb{N}}$ are the eigenvalues of A_t , then

$$\|A_t^{-1}\|^{-1} = \min_{k \geq 1} \alpha_{t,k} \geq \underline{\alpha}_t, \quad \|I - \gamma_t A_t\| = \max_{k \geq 1} (1 - \gamma_t \alpha_{t,k}) \geq 0,$$

where we use that, for all $k \in \mathbb{N}$, $\gamma_t \alpha_{t,k} \leq \gamma_t \bar{\alpha}_t \leq 1$.

But $\min_{k \geq 1} \alpha_{t,k} \geq \underline{\alpha}_t$ implies $\max_{k \geq 1} (1 - \gamma_t \alpha_{t,k}) \leq 1 - \gamma_t \underline{\alpha}_t$, thus (A).

The last claim follows from the inequality

$$\prod_{i=j}^t \left(1 - \frac{c}{i + t_0} \right) \leq \exp \left(- \sum_{i=j}^t \frac{c}{i + t_0} \right) \leq \exp \left(-c \log \left(\frac{\bar{t} + 1}{j + t_0} \right) \right) = \left(\frac{j + t_0}{\bar{t} + 1} \right)^c.$$

□

Proof of Theorem 3.5. First, $\gamma_t \bar{\alpha}_t \rightarrow 0$ implies that there exists $j_0 \in \mathbb{N}$ such that $\gamma_t \bar{\alpha}_t \leq 1$ for all $t \geq j_0$. Hence Lemma 3.8 (B) applies, so that

$$(C-1) \quad \|\Pi_j^t\| \leq \prod_{i=j}^t (1 - \gamma_i \underline{\alpha}_i).$$

Let us use the reversed martingale decomposition of $r_.$, from times j_0 to t :

$$\|r_t\| \leq \mathcal{E}_{init}(t) + \mathcal{E}_{samp}(t) + \mathcal{E}_{drift}(t),$$

where

$$\mathcal{E}_{init}(t) := \|\Pi_{j_0+1}^t r_{j_0}\|, \quad \mathcal{E}_{samp}(t) := \left\| \sum_{j=j_0+1}^t \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j) \right\|, \quad \mathcal{E}_{drift}(t) := \left\| \sum_{j=j_0+1}^t \Pi_j^t \Delta_j \right\|.$$

Now, by (C-1),

$$\mathbb{E}(\mathcal{E}_{init}(t)^2) \leq \exp \left(-2 \sum_{i=j_0+1}^t \gamma_i \underline{\alpha}_i \right) \mathbb{E}(\|r_{j_0}\|^2) \xrightarrow{t \rightarrow \infty} 0$$

since $\sum_t \gamma_t \underline{\alpha}_t = \infty$, and

$$\mathcal{E}_{drift}(t) \leq \sum_{j=j_0+1}^t \|\Delta_j\| \prod_{i=j}^n (1 - \gamma_i \underline{\alpha}_i) \xrightarrow{t \rightarrow \infty} 0$$

by assumption (C'). Now consider the sample error. Using the independence of $(z_t)_{t \in \mathbb{N}}$,

$$\begin{aligned} \mathbb{E}(\mathcal{E}_{samp}(t)^2) &= \mathbb{E} \left\| \sum_{j=j_0+1}^t \gamma_j \Pi_{j+1}^t (A_j \bar{w}_j - b_j) \right\|^2 = \sum_{j=j_0+1}^t \gamma_j^2 \mathbb{E} \|\Pi_{j+1}^t (A_j \bar{w}_j - b_j)\|^2 \\ &\leq C \sum_{j=j_0+1}^t \gamma_j^2 \prod_{i=j+1}^t (1 - \gamma_i \underline{\alpha}_i)^2, \end{aligned}$$

where $C := \sup_{t \in \mathbb{N}} \mathbb{E} \|A_t \bar{w}_t - b_t\|^2 < \infty$ by assumption. This completes the proof, using (C'). \square

Proof of Lemma 3.6. Let $\epsilon > 0$. The assumptions $\limsup_{t \rightarrow \infty} a_t/b_t = 0$ and $b_t \rightarrow_{t \rightarrow \infty} 0$ imply that there exists $N \in \mathbb{N}$ such that $a_t \leq \epsilon b_t/2$ and $b_t \leq 1$ for all $t > N$. On the other hand, $\sum_{t \in \mathbb{N}} b_t = \infty$ implies that there exists $N_1 \in \mathbb{N}$ such that, for all $n \geq N_1$,

$$\sum_{k=1}^N a_k \prod_{i=k+1}^n (1 - b_i) < \frac{\epsilon}{2}.$$

Now

$$\sum_{k=N+1}^n a_k \prod_{i=k+1}^n (1 - b_i) \leq \frac{\epsilon}{2} \sum_{k=N+1}^n b_k \prod_{i=k+1}^n (1 - b_i),$$

and we can write the right-hand side of this last inequality as a telescopic sum, i.e.

$$\sum_{k=N+1}^n b_k \prod_{i=k+1}^n (1 - b_i) = \sum_{k=N+1}^n [1 - (1 - b_k)] \prod_{i=k+1}^n (1 - b_i) = 1 - \prod_{i=N+1}^n (1 - b_i) \leq 1,$$

which enables us to conclude. \square

REFERENCES

BERLINET, A. and C. THOMAS-AGNAN (2004). *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer Academic Publishers.

CAPONNETTO, A. and E. DE VITO (2007). Optimal rates for regularized least squares algorithm. *Found. Comput. Math.* 7(3), 331–368.

CARMELI, C., E. DEVITO, and A. TOIGO (2005). Reproducing kernel hilbert spaces and mercer theorem. [arXiv:math/0504071v1 \[math.FA\]](https://arxiv.org/abs/math/0504071v1).

CUCKER, F. and S. SMALE (2002). On the mathematical foundations of learning. *Bull. of the Amer. Math. Soc.* 29(1), 1–49.

DAUBECHIES, I. (1992). *Ten Lectures on Wavelets*. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM).

DUFLO, M. (1996). *Algorithmes Stochastiques*. Berlin, Heidelberg: Springer-Verlag.

EFRON, B., JOHNSTONE, I., HASTIE, T., and TIBSHIRANI, R. (2004). *Least angle regression*. *Annals of Statistics* 32(2), 407–451.

ENGL, H. W., HANKE, M. and NEUBAUER, A. (1996). *Regularization of Inverse Problems. Mathematics and its Applications*, 375. Kluwer Academic Publishers Group.

T. EVGENIOU, M. PONTIL, and T. POGGIO (2000). *Regularized networks and support vector machines*. *Adv. Comput. Math.* 13, 1–50.

GYÖRFI, L., M. KOHLER, A. KRZYŻAK, and H. WALK (2002). *A Distribution-Free Theory of Nonparametric Regression*. Springer-Verlag, New York.

HALMOS, R. P. and V. S. SUNDER (1978). *Bounded Integral Operators in L^2 Spaces*. Vol. 96 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (Results in Mathematics and Related Areas)*. Berlin: Springer-Verlag.

J. KIEFER and J. WOLFOWITZ (1952). Stochastic estimation of the maximum of a regression function. *Ann. Math. Statist.* 23, 462–466.

KIVINEN, J., A. J. SMOLA, and R. C. WILLIAMSON (2004). Online learning with kernels. *IEEE Transactions on Signal Processing* 52(8), 2165–2176.

KUSHNER, H. J. and G. G. YIN (2003). *Stochastic Approximations and Recursive Algorithms and Applications*. Berlin, Heidelberg: Springer-Verlag.

LOÉVE, M. (1948). Fonctions aléatoires du second ordre. In P. Lévy (Ed.), *Processus Stochastiques et Mouvement Brownien*, Paris. Gauthier-Villars.

NEVEU, J. (1975). *Discrete-Parameter Martingales*. North-Holland Publishing Company.

PARZEN, E. (1961). An approach to time series analysis. *The Annals of Mathematical Statistics* 32(4), 951–989.

PINELIS, I. (1994). Optimum bounds for the distributions of martingales in banach spaces. *The Annals of Probability* 22(4), 1679–1706.

ROBBINS, H. and S. MONRO (1951). A stochastic approximation method. *The Annals of Mathematical Statistics* 22(3), 400–407.

ROSSET, S. and J. ZHU (2007). Piecewise linear regularized solution paths. *The Annals of Statistics* 35(3), 1012–1030.

SMALE, S. and Y. YAO (2006). Online learning algorithms. *Foundation of Computational Mathematics* 6(2), 145–170.

SMALE, S. and D.-X. ZHOU (2004). Shannon sampling and function reconstruction from point values. *Bull. of the Amer. Math. Soc.* 41(3), 279–305.

SMALE, S. and D.-X. ZHOU (2007). Learning theory estimates via integral operators and their approximations. *Constr. Approx.* 26(2), 153–172.

TARRÈS, P. and Y. YAO (2008). Learning as stochastic approximation of regularisation paths. *Report 30/2008, Learning Theory and Approximation workshop, Mathematisches Forschungsinstitut Oberwolfach*.

TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc. B.* 58(1), 267–288.

WAHBA, G. (1990). *Spline Models for Observational Data*. SIAM. CBMS-NSF Regional Conference Series in Applied Mathematics, 59.

WILLIAMS, D. (1991). *Probability with Martingales*. Cambridge University Press.

YAO, Y. (2006). *A Dynamic Theory of Learning – Online Learning and Stochastic Algorithms in Reproducing Kernel Hilbert Spaces*. PhD Dissertation, University of California at Berkeley, published by Verlag Dr. Muller, ISBN: 978-3-639-09390-2.

YAO, Y. (2010). On complexity issue of online learning algorithms. *IEEE Transactions on Information Theory* 56(12), 6470–6481.

PIERRE TARRÈS, MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, 24-29 ST GILES', OXFORD OX1 3LB, U.K.

E-mail address: tarres@maths.ox.ac.uk

YUAN YAO, SCHOOL OF MATHEMATICAL SCIENCES, PEKING UNIVERSITY, BEIJING, CHINA 100871.

E-mail address: yuany@math.pku.edu.cn